

# Using SDN-aware Service Assurance to Bring Open and Deterministic Service Management to NFV

With a virtualized Service Assurance Manager (vSAM) solution, KPIs are abstracted and open, enabling service providers to improve resource usage and quality of experience across the entire SDN/NFV architecture.



## Introduction

Service assurance accurately measures and reports on the infrastructure (network and platform) key performance indicators (KPIs) that may affect a specific service. Applying concepts from software-defined networking (SDN) and network functions virtualization (NFV), this paper outlines a service assurance approach that enables more open and deterministic service deployment and resource usage. Whilst the approach is generic and NFV-application agnostic, it is discussed in the context of a Gi-LAN example to show the value it may bring to service providers. The approach has similar benefits for vCPE, vIMS, vEPC, and other virtualized applications.

The paper also explains the concept of a Virtual Service Assurance Manager (vSAM), which gives network operators real-time visibility into SDN/NFV-based equipment performance and loading across the Wide Area Network (WAN). This information can be used to perform optimized load balancing in large networks in order to improve equipment utilization and quality of experience (QoE). The vSAM marries service assurance capabilities at two levels: equipment platform and WAN. An equipment platform solution is provided by the Intel® Service Assurance Administrator (Intel® SAA), and WAN service assurance solutions are delivered by third-party vendors.

Why Is Service Assurance Needed in an NFV Environment?

The main concerns service providers have about NFV deployment are very closely linked to the determinism and service assurance capability of the NFV infrastructure (NFVI). This is confirmed by the Heavy Reading\* NFV Market Tracker shown in Figure 1, where five of the top eight challenges are directly impacted by the service assurance capabilities of the SDN/NFV system.

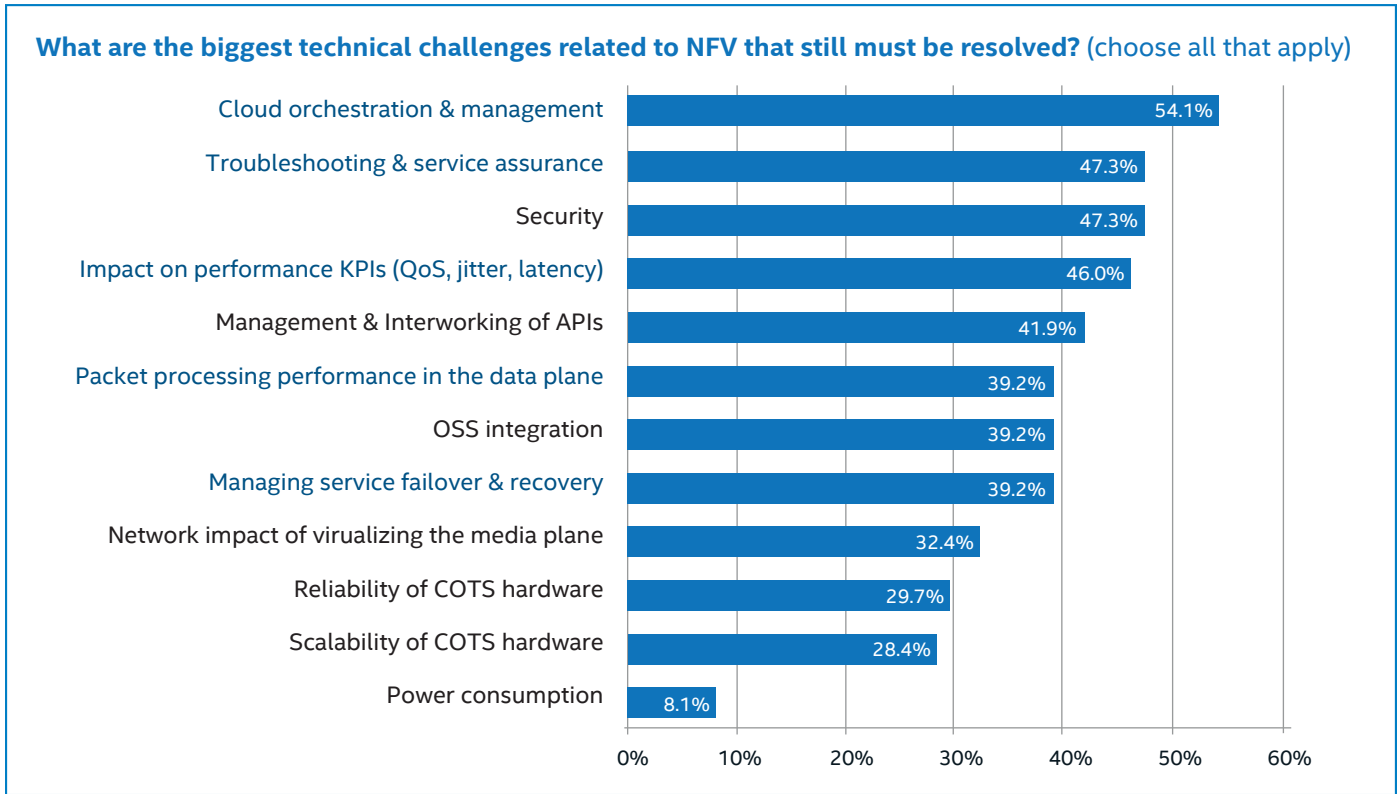


Figure 1. Operator Concerns with NFV Source: Heavy Reading\*, Q4 2014

At the same time, there are several network requirements that must be considered, including:

- **Telco workloads are more demanding than typical cloud workloads.** Whilst a web-based application may survive quite 'long' outages and variances in resource or network performance, many telecommunications protocols are architected for very strict performance windows, and indeed, may even take undesired, unilateral action, such as link or node protection schemes, should the infrastructure not serve the protocol within the required window.
- **Service providers are mandated to deliver real-time, always-on services, like 911.** This requires a service provider to employ a service management layer that is fully aware of all the resource KPIs in real time and with full determinism. Contravening regulatory compliancy can have a significant impact on the service provider's ability to operate.
- **Appliance vendors are measured and contracted by the uptime of their equipment.** Vendors failing to meet the uptime SLAs can face commercial penalties. Consequently, all vendors in the existing appliance-centric model integrate redundancy and uptime mechanisms into the vertical solution. In the virtualized environment, with multiple vendors providing functions in a shared environment, the bar on service assurance tools is even higher. NFV service assurance tools need to monitor not only that every function is living up to its promised KPI, but also that the shared infrastructure is providing adequate resources in real time.

Appliance vendors have traditionally built in compliance towards KPIs and service assurance capability by taking advantage of a very tight coupling between the appliance hardware, software, and the associated vendor element management system (EMS) and network management system (NMS). This was necessary as only the vendors (the appliance and software designers) were deeply aware of the scaling limits and performance KPIs of their hardware systems, and each appliance was different in terms of architecture. With NFV, the need for tight coupling between software and the underlying hardware resources disappears. In fact, a tight coupling presents an obstacle for modularity and diversity in the SDN/NFV ecosystem.

With NFVI, service assurance functions have to be abstracted to enable the tight vendor Operational Support System (OSS)-VNF linkages to be broken and full NFV modularity to be achieved. This allows services to be deployed between different vendor VNFs from an upper layer controller with full determinism and visibility. It is Intel's view that the NFVI has to present an open standard, northbound interface to the hypervisor and VNF layers. In addition, NFVI needs open standards to support east-west traffic going across logical links into other VNFs in the same or remote NFVI Point of Presence (PoP).

### Gi-LAN Basics

A Gi-LAN is the part of the network that connects the mobile network (left side of Figure 2) to data networks, such as the Internet and operator cloud services. The term "Gi" refers to the 3GPP reference point between the Gateway GPRS Support Node (GGSN) and a packet data network (PDN). The Gi-LAN can have upwards of fifty assorted appliances running applications such as deep packet inspection (DPI), firewall, URL filtering, DMZ, Network Address Translation (NAT), and video and web optimizers, etc.

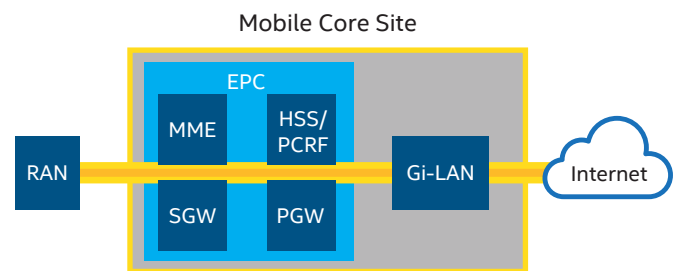


Figure 2. Gi-LAN Placement in the Network

### Operator Challenges

Mobile core sites cater to millions of subscribers and are highly complex with respect to telecom architecture. These sites include 3G, Evolved Packet Core (EPC), and Gi-LAN (see sidebar) functions with the associated management and policy functions. There are many appliances from different vendors, many protocols to support, and stringent SLAs. Ironically, this is also the part of the network that needs to be most agile and programmable in terms of service delivery and subscriber QoE.

A primary SDN/NFV use case for mobile network operators (MNOs) is service chaining in the Gi-LAN. This can provide a wide range of benefits, including:

- **On-Site Resource Usage Optimization:** This capability ensures that all the traffic does not take a serial path through the network and that different types of traffic traverse only the appropriate functions. For example, smartphone traffic from a sports event may be sent to the video optimizer pool while data from a fitness tracker (attached to the phone via a personal area network) could be steered towards a specific analytics engine.

- **Service Velocity.** Today's appliance-built Gi-LANs are very complex and rigid. Normal operational activities, like software upgrades, new bundle introduction, new policy introductions, or appliance upgrades, can take several weeks and incur considerable OpEx. SDN and NFV architectures with 3GPP-integrated control planes are expected to reduce this dramatically.
- **Improved Vendor Modularity.** Appliance-based Gi-LANs make it extremely difficult to insert or remove incumbent vendors. With service chaining and NFV, vendor VNFs may be inserted, tested, deployed, or removed with much greater ease.

It should be noted that the industry has several competing standards initiatives to implement service chaining. There is still debate on tunneling schemes, service contexts, on the granularity of integration between SDN and 3GPP policy planes in the service function classifier (SFC), and on the traffic detection function (TDF). Service chaining will have to support existing appliances and VNFs, and ensure the method used for this scales and does not introduce further complexity. Proof of concepts and innovation to address these issues are on-going.

**Solution Outline**

To provide operators an open view of resources on a network-wide basis, Intel is introducing the concept of a vSAM, which combines the platform assurance capabilities delivered by the Intel® Service Assurance Administrator (Intel® SAA) with WAN service assurance capabilities offered by third parties. It feeds summary results into a northbound service chaining system and third-party OSS functions. There is one vSAM per NFVI PoP.

Intel SAA provides essential software tools for OpenStack\* cloud services and infrastructure management. The product enables the creation of a software-defined infrastructure with enhanced service level objectives. Intel SAA runs on the host OS and monitors all virtual machine (VM) activity in terms of CPU, memory, and I/O usage. The key features and benefits of Intel SAA are listed in Table 1.

FEATURES	BENEFITS
Rich SLA and NFVI alignment	<ul style="list-style-type: none"> <li>• Target machine instances to run only on trust-attested compute nodes</li> <li>• Specify performance quota with Service Compute Unit enabling higher density of VMs per node</li> </ul>
Automated provisioning with intelligent machine placement	<ul style="list-style-type: none"> <li>• Automate placement of VMs on trust-attested nodes</li> <li>• Automate placement of VMs to avoid performance issues</li> </ul>
OpenStack* health monitoring: efficient administration	<ul style="list-style-type: none"> <li>• Monitor critical OpenStack* components</li> <li>• Monitor compute nodes in-depth</li> </ul>
Analysis and remediation engine: probable root cause analysis	<ul style="list-style-type: none"> <li>• Collect log file data</li> <li>• Detect anomalies</li> <li>• Detect SLA violations</li> </ul>
Capacity and usage statistics for insightful planning	<ul style="list-style-type: none"> <li>• Compute node resource capacity, capability, and consumption</li> <li>• Monitor VM capacity, capability, and consumption</li> <li>• Automate intelligent machine instance provisioning</li> </ul>

**Table 1.** Keys Features and Benefits of Intel® Service Assurance Administrator (Intel® SAA)

**ContexNet SDN Fabric Solution**

ContexNet from ContexStream is a distributed overlay SDN fabric solution designed for NFV. The solution maps in new, real-time flows, which can be identified in various granularities, to available function resources: physical or virtual. The solution comprises several components:

- the **switch** is an OpenFlow\* entity in the traffic path
- the **controller** generates the rules
- the **mapping service** provides the distributed controller entities a centralized view of resource location/identity, utilization etc.

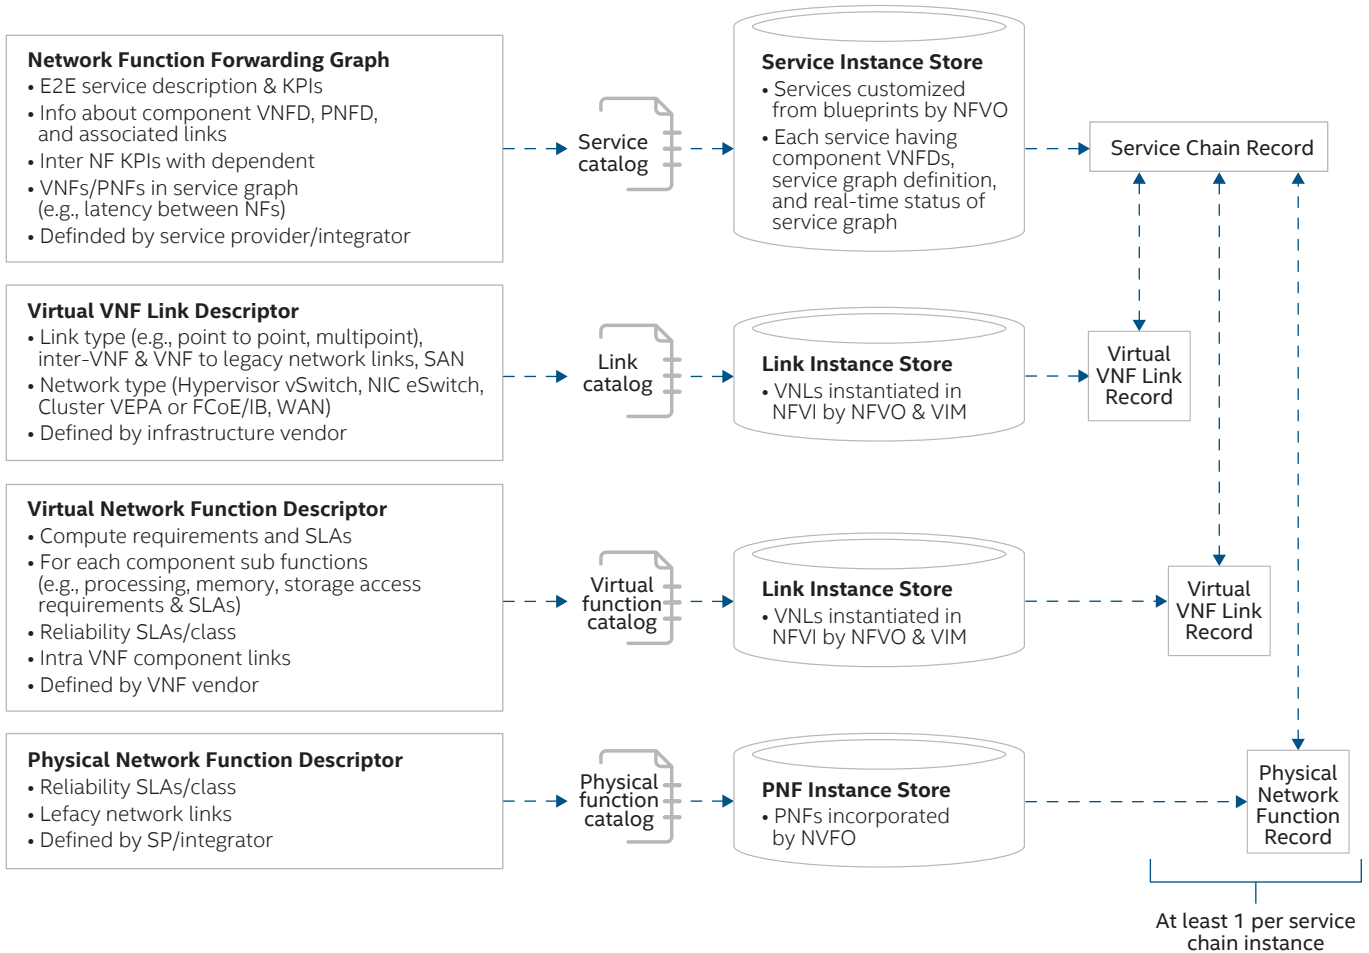
The mapping service is based on a distributed database, and the controller interfaces to it using a LISP (IETF standard) framework that has been extended to a PUB/SUB paradigm. When a new traffic flow starts, the controller consults the mapping service for the “currently best available” instance of the function needed within the cloud fabric and then installs rules for the traffic flow to be directed to that function instance in real time. ContexNet, along with Intel SAA, will add to the available information to enable the controller to make service-aware decisions.

**vSAM Functions**

By merging equipment platform and WAN service assurance capabilities, the vSAM is able to:

- Measure real-time load on all local VNFs (CPU, memory, I/O, and session count)
- Measure inter-NFVI-PoP network conditions (bandwidth, packet loss, delay, and delay variation)

- Expose the information to the service function chaining (SFC) for constraint-based service chaining either on-site (locally) or over the WAN for failure and overload scenarios. This will be achieved by populating the link and descriptors on the left side of the ETSI framework depicted in Figure 3, and presenting this information holistically for third-party systems to utilize as required.

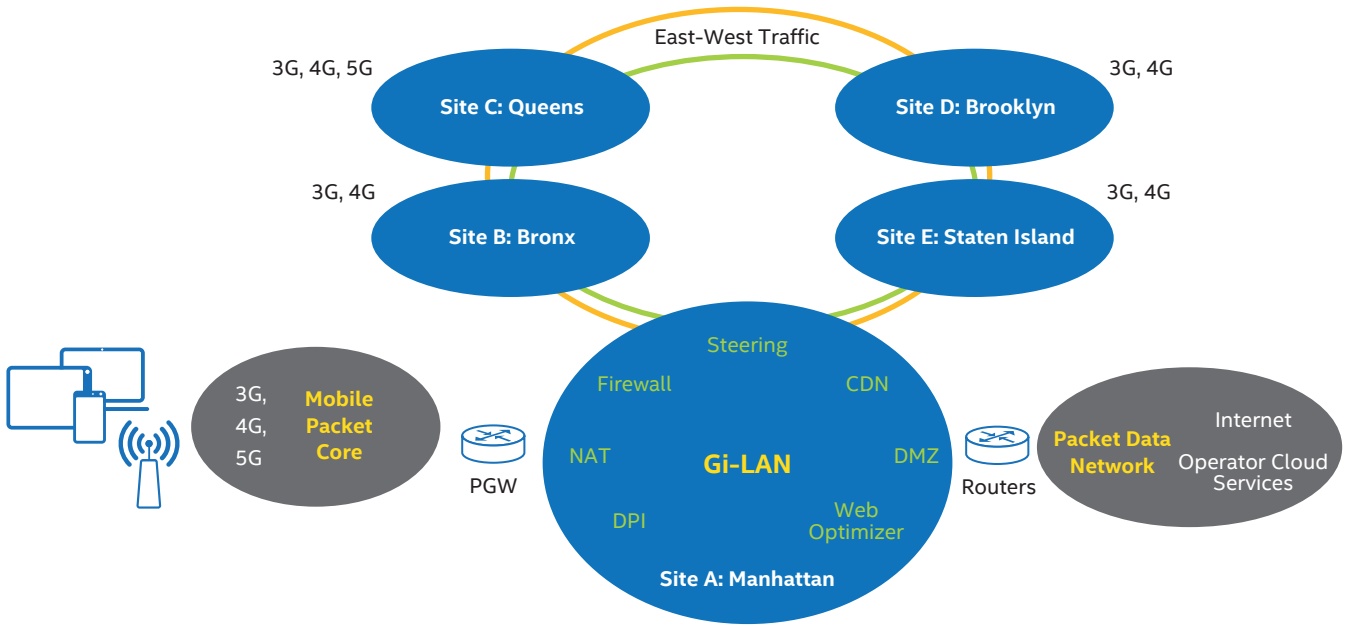


**Figure 3.** ETSI Service Chaining Information Model with Service Assurance Descriptor KPIs

By combining the platform and network KPIs, vSAM presents a real-time network view of all NFVI resources to the service chaining architecture.

**Inter-Site Implications**

There is also a wider context to service chaining. In a large, densely populated urban area (conurbation), there could be multiple Gi-LAN sites. In terms of resiliency, the sites are dimensioned to cope with a single failure at any adjacent site. Figure 4 shows a hypothetical example of five Gi-LAN sites serving the boroughs of New York City, such that if any of the Gi-LANs fails, the other four sites should be able to handle the peak subscriber load during the busy hour.



**Figure 4.** Site-to Site Gi-LAN Architecture

At a network level, this architecture presents some challenges:

- 1. Coarse Failover:** These sites are very complex and static in terms of configuration, and the protection schemes between sites are very coarse. For example, if a packet data network gateway (PGW) or a carrier-grade NAT fails, the Multi-Protocol Label Switching (MPLS) Provider Edges (PEs) on the Gi-LAN site typically reroute all the traffic from one mobile service core to another. PEs do this without knowledge of how loaded the remote CG-NATs or PGWs are; and as a result, traffic could be routed to an overloaded site and service levels could significantly degrade.
- 2. Site Overbuild:** This lack of awareness presents a large CapEx problem for service providers, in that the only way to provide determinism for subscriber QoE under failure conditions is to massively overbuild these sites. The overbuild situation becomes even worse with non-symmetrical adjacent sites, whereby a small-core site is required to handle the load of an adjacent, large-core site in the event of failure.
- 3. Site Overload:** In addition, today's Gi-LAN sites are not good at dealing with unexpected local overload. Although all core networks are deliberately oversubscribed, extraordinary events in a particular city can unexpectedly overload the local Gi-LAN. This can be prevented with methods that dynamically load share services to other resources during an overload, thus avoiding black-holed traffic and frustrated subscribers.
- 4. Vendor Modularity:** Current appliance architectures make the introduction of new 3GPP technologies very expensive. In the example shown in Figure 4, 3G and 4G are served by all five sites so the operator needs to overbuild all sites by 25 percent to cope with a single-site failure assuming equal-sized sites. Since 5G is served by only two sites, the MNO has to overbuild the Gi-LAN by 100 percent for 5G services. This situation will be repeated with upcoming 3GPP releases as the industry moves to new technologies.

### Example Solutions

Two potential use cases are depicted in Figure 5. The example starts with the SDN infrastructure deploying two service chains on the local Gi-LAN in Manhattan (bottom of figure). One service chain directs the streaming traffic (purple line) towards the local content delivery network (CDN) and video optimizer (VO), and another service chain (pink line) sends voice traffic to the local session border controller (SBC) for the VoLTE client.

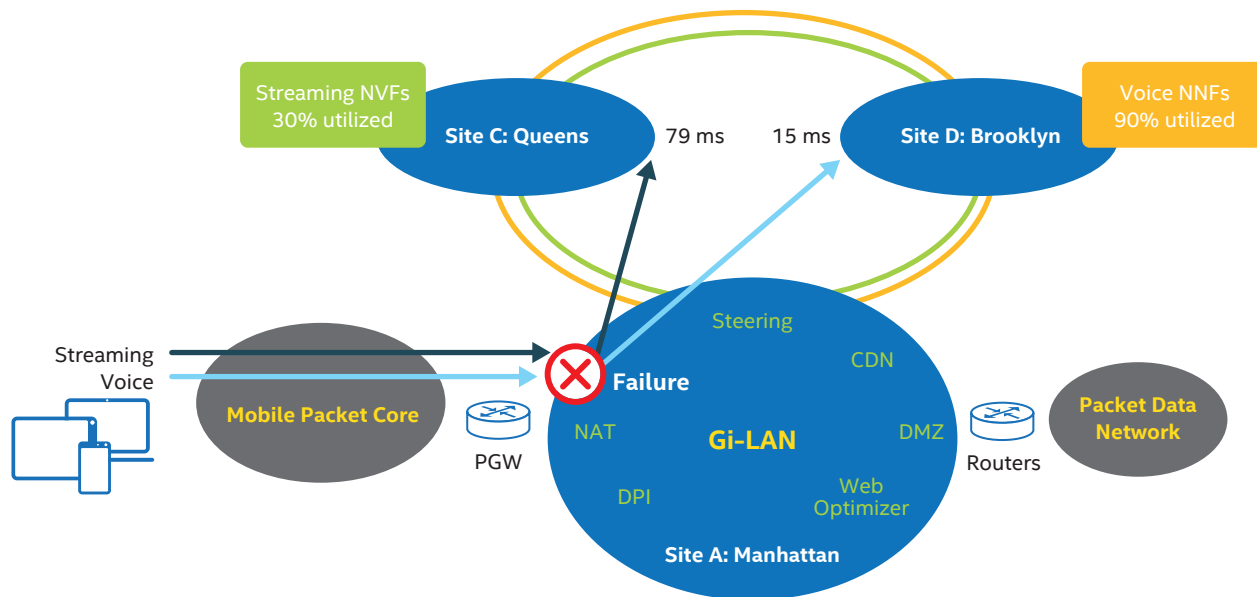


Figure 5. Gi-LAN Inter-NVI PoP Traffic Steering Use Cases

Continuing with this theoretical example, two failures occur in the local Gi-LAN in Manhattan.

1. The local VO and CDN in Manhattan become unavailable. Since the streaming application is not delay sensitive, the SFC, working in conjunction with the vSAM, redirects the service chain towards the Gi-LAN in Queens (left side), which is lightly loaded, but on a slow virtual link. Thus, the service is not severely impacted by the local failures. This scenario applies equally to situations where the local VO or CDN are heavily loaded and unable to service new flows.

2. The local SBC handling the voice traffic becomes overloaded or unavailable. Here the SFC decides to chain the traffic to the Gi-LAN in Brooklyn (right side) for SBC functionality, even though Brooklyn is heavily loaded. This is because the voice application is delay sensitive, making the faster virtual link preferable (15 ms versus 70 ms latency).

This example demonstrates how to implement more deterministic load sharing and granular failover of virtualized Gi-LAN resources across the WAN in real time. The solution offers sophisticated service chaining to address service provider OpEx concerns and intelligent resource sharing between NFVI-PoPs to help reduce CapEx.

## Conclusion

Service providers have led the industry in terms of SDN/NFV investigation, proofs of concepts, and now, adoption. The benefits of the SDN/NFV architectural model are clear to the service provider community. However, key concerns remain around the SDN/NFV-enabled network centre with respect to the level of determinism, management, and resilience that can be achieved in an open and modular way.

The specific application benefits outlined in this paper were regarding Gi-LAN, but the approach is valid for all applications that need to share resources intra-Pop or inter-PoP. As such vEPC, vIMS, and other applications could also greatly benefit from an open service assurance fabric for NFV. Integral to this solution is Intel SAA, which provides a way to collect service assurance information from equipment platforms.

By implementing a vSAM horizontal architecture, whereby service assurance KPIs are abstracted and open, service providers can remove tight dependencies between upper layer service management systems and NFVI. This ensures NFVI presents an open interface from a platform and network point of view, which helps address operator challenges around vendor modularity, on-site resource usage optimization, and service velocity across the whole SDN/NFV architecture.

For more information about Intel® solutions for communications infrastructure, visit [www.intel.com/go/commsinfrastructure](http://www.intel.com/go/commsinfrastructure)

For more information about NFV- and SDN-based solutions, visit <https://networkbuilders.intel.com>

Acronyms	Term
3GPP	3rd Generation Partnership Project
API	Application Programming Interface
DPDK	Data Plane Development Kit
DPI	Deep Packet Inspection
E2E	End to End
EMS	Element Management System
EPC	Evolved Packet Core
ETSI	European Telecommunications Standards Institute
Gi-LAN	Gateway-Internet Local Area Network
GTP	Gateway Tunneling Protocol
IETF	Internet Engineering Task Force
IMEI	International Mobile Equipment Identity
IP	Internet Protocol
LTE	Long Term Evolution
MANO	Management and Orchestration
MPLS	Multiprotocol Label Switching
NETCONF	Network Configuration Protocol
NFV	Network Function Virtualisation
NFVI	Network Function Virtualisation Infrastructure
NMS	Network Management System
NFV	Network Function Virtualization
NSH	Network Service Header
TDF	Traffic Detection Function
TWAMP	Two-Way Active Measurement Protocol
VNFFG	Virtual Network Function Forwarding Graph
VNF	Virtual Network Function
VXLAN	Virtual Extensible LAN

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request. Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order. Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting Intel's Web site at [www.intel.com](http://www.intel.com).

Copyright © 2015 Intel Corporation. All rights reserved. Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.

