

Executive Summary

Virtualization of different segments and functions of wireline and wireless network infrastructure will help ensure that service providers' business key performance indicators (KPI) remain at healthy levels. Today, some of these parameters are under pressure with the top line growth limited due to slowing of new customer acquisition, competitive price pressures, and a constantly increasing set of new services from OTT players, while traffic volumes and operational costs continue to rise. Software defined networking (SDN) and network function virtualization (NFV) help circumvent stringent limitations in growth, operation, and new service offerings. The immediate impact of SDN/NFV based networks is network infrastructure optimization and operational simplification, leading to lower total cost of ownership (TCO). TCO savings notwithstanding, it is the revenue generation possibilities that are expected to have a deeper longer term positive impact on an operator's business and in ACG's opinion must be given more weight than reduction of TCO. With an SDN/NFV based infrastructure, operators will be able to scale their service offering to levels that were simply unfeasible to achieve with purpose-built network components. The software-centric nature of the infrastructure provides operators an unprecedented ability to implement and deploy a diverse portfolio of services to maximize their revenues and while reducing the inherent risk of failure associated with a new service offering. Hewlett Packard Enterprise has emphasized and addressed these with its HPE ContextNet solution. This is a carrier-grade SDN/NFV solution for operators to migrate from present mode of operation (PMO) of their network to implement, and deploy a virtualize mode of operation (VMO). The solution's various components enable service providers to manage/provision/orchestrate P+V VNFs in Gi LAN, resulting in faster time-to-market in deployment of new services.

ACG Research conducted an analysis of a typical scenario that a service provider faces when migrating to a VMO based Gi LAN network to expand network scale. The scenario assumes an existing Gi LAN infrastructure and the choice to continue with PMO versus a VMO based deployment. The study found a reduction of TCO level, but more importantly a substantial gain in revenue volume based on only two new services (audio/video streaming and fee-based Apps) over five years as a result of faster time-to-market for a "new service" creation (32 percent of the time it would take to deploy the same services relative to the PMO solution). The study also found that per increase in scaling of new services for VMO, an additional cumulative revenue of \$27 million can be realized for every 50,000 subscribers that sign up with SP's increased service portfolio. For every 3 percent margin of error in traffic volume forecast, close to \$1.3 million was wasted per over-dimensioning the network infrastructure.

KEY FINDINGS

HPE ContextNet carrier-grade SDN/NFV fabric enables a smooth migration from purpose-built Gi LAN infrastructure to a virtualized mode of operation. Over five years:

- Deploy a new service in 1/3 of the time it would take the PMO to deploy the same service
- For two services, analyzed, an additional \$54.28 million in revenue is gained per faster TTM
- \$27 million additional cumulative revenue, over a 5 year period, for every 50,000 subscribers gained via service agility
- Waste of \$1.3 million for a 3% error in traffic volume forecast for PMO per over-dimensioning
- TCO savings of 114%

Introduction

The Gi LAN is a critical segment of mobile (and converged) network infrastructure to deploy services demanded by today's mobile centric society. It is in this segment of the network core where operators can utilize to offer new services and deploy new business models. Typically for vanilla mobile Internet access network security (Firewall) and CG-NAT are essential functions in the Gi-LAN, but this is the segment of the network where operators can deploy (experiment) with new business models and services. For e.g. operators can add new value-added services by introducing

1. Optimization and caching of video-content,
2. Customized service levels via policies per subscriber with policy enforcement,
3. Analytics to monitor and take action per network performance and
4. Traffic optimization.

More importantly, the Gi LAN services enable service providers (SP) to offer innovative business models and new services, for example, premium video/audio streaming, partnership with popular social media Apps for subsidizing subscriber costs in non-3GPP based access, for example, Wi-Fi. A typical serially connected Gi LAN network is shown in Figure 1.




Figure 1. Typical Gi LAN Architecture and Placement in Packet Core

The current architecture utilizes a series of appliances that each offers a distinct service that comprise the Gi LAN. The services can and have been integrated in one chassis, running on different cards. There are several inherent drawbacks with this architecture:

- Introduction of a new appliance or scaling an existing appliance, for example, add more capacity to content caching, is rigid, time-consuming and may not be deemed to be economical.

- The architecture is not conducive to deliver new services within a time frame that could prove to be lucrative for the service provider as it requires long cycle time frames to provision, test and deliver a successful new service.
- As traffic increases, the number of load balancers may also increase, in step, adding more complexity to integration and management of Gi LAN.
- Traffic runs through each service in a serial fashion whether necessary, potentially increasing latency.

HPE ContextNet (Figure 2), a carrier-grade, distributed SDN Fabric takes advantage of virtualization enabling service providers to:

- Automatically introduce and withdraw a new appliance or service to capture market share and subscription “stickiness” for higher revenue levels.
- Benefit from a simpler network architecture to reduce the need serial traffic flow (less latency) and eliminate the need for standalone deployment of certain Gi LAN functions, for example, load-balancers, analytics, policy enforcement to help further contain TCO.
- Operate a programmable network to run “service-chaining,” via third-party VNFs of choice with subscriber awareness at a most granular service level to a subscriber to implement innovative business models for additional revenue streams.
- Manage the network for P+V for a smooth migration from PMO to VMO with elasticity for on-demand resource allocation to avoid under/over-provisioning, which can lead to lost opportunities or lower return on asset (ROA).




Figure 2. HPE ContextNet for Deployment of Virtualized Gi LAN

HPE ContexNet is composed of three main components: ContexSwitch, ContexControl, and ContexMap, Figure 3, which shows the interfaces and interactions of all the components of the SDN fabric.

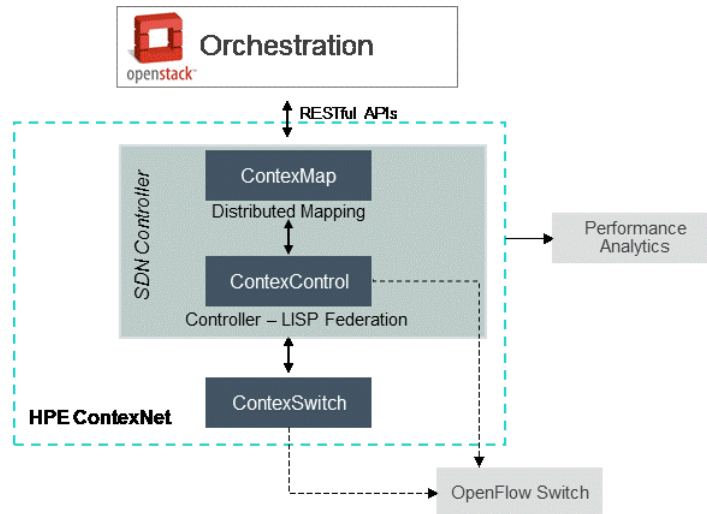


Figure 3. Interfaces and Interactions of Components of SDN Fabric

Table 1 lists the high-level features and their respective benefit for each HPE ContexNet component.

Module	Functions/Features	Benefits
ContexMap	ContexMap is a distributed mapping subsystem that enables subscriber lookup and network function information from any HPE ContexNet node.	Efficient use of network resources and tools to implement innovative service/business models in a distributed operator environment.
ContexControl	ContexControl is a carrier-class SDN Controller based on a federated control; efficiently load balances all traffic on all resources. It is service, subscriber and application aware and only sends traffic to a function if needed per subscriber flow. Controllers are physically distributed but logically centralized through ContexMap.	Introduction of a programmable network architecture and lower overall automation cost
ContexSwitch	ContexSwitch is the high-performance software OpenFlow switch that augments the capabilities of hardware OpenFlow switches. Support for Overlay Network in the network (OpenFlow switches) not	Support for packet flows in software delivers a high performance network and flexible network that can potentially enable

	on the host to maintain traffic treatment as required for scalability and reliability.	programmability with lower cost of operations and new function introduction.
--	--	--

Table 1. HPE ContexNet Features and Benefits

HPE ContexNet Economic Advantages

HPE ContexNet tools offer service agility for the deployment of new services for a faster time-to-market (TTM) relative to the PMO based infrastructure. This is essential for service providers to increase subscriber take rate and revenue streams and keep their competitive advantage. PMO solutions simply cannot offer the rapidity and scale of services that an operator would require for introducing a new service to its target market. PMO based service introduction has proved to be expensive and not necessarily always successful, making them somewhat of a risky venture. In addition, the service delivery time frames have long cycles potentially limiting the number of services that an operator can offer within a period of time, for example, for a calendar a year, a maximum of two to three new services depending on availability of internal resources.

The PMO service delivery timeframes drawbacks have been solved by HPE ContexNet:

- Rapid implementation/experimentation and deployment of a service:
 - The service provider is not bound by the number of services that it can offer within a similar period. The scale can now be in 10s or can even be 100s of services. These can be quickly agreed upon among the service provider’s decision makers and influencers, implemented, market tested and deployed with a fraction of the cost that a PMO based infrastructure would have to incur. With HPE ContexNet the service provider has the luxury of “fast-failure” for a service without incurring the level of the cost that is faced by a PMO. Today operators have to compete with OTT providers (e.g., Dropbox, Whatsapp, Uber) and cannot afford the high cost and long deployment cycles associated with PMO
 - A comparison in service delivery timeframe can be shown by a new index, ACG calls, Relative Agility (RA). RA measures the agility of a specific solution (S_1) that can create a new resource versus its competing solution (S_2), capable of creating the same resource. S_1 is the “dev ops”¹ model of innovation and S_2 is the typical model used by operators with physical proprietary infrastructure. Therefore, the RA for S_1 vs. S_2 gives us which solution is more agile in introduction of a new resource (in this case, we assume the “resource” is a new service). Assuming that agility of a virtualized solution is X, as in X days, say this is S_v (for a virtualized solution), then the relative agility of S_p , a PMO based solution, that is, non-virtualized, to S_v , for introducing the same resource would be:

¹ <https://www.ietf.org/id/draft-unify-nfvrg-devops-03.txt>

$R.A = (MTRI \text{ for } S_v)/(MTRI \text{ for } S_p)$, where MTRI = Mean Time to Resource Introduction. The time variable could be days/weeks/months

For example, if it takes S_p five days to introduce the same service that S_v is able to complete in two days, its relative agility for that service is $2/5$ or 40 percent. In our study for HPE ContextNet versus PMO, and assuming no changes beyond the Gi-LAN (e.g. to OSS/BSS processes), this figure is calculated to be $R.A_{PMO} = 44/139 = 32$ percent. In other words we found that, with HPE ContextNet a service can be introduced in about one-third of the time that it would take to introduce the same service via PMO. HPE ContextNet provides real-time, on-demand availability of a service at scale. Relative Agility can be used as a rule of thumb in factoring other variables such TTM for a new service. The smaller the R.A., the higher opportunity cost for the service provider using the slower solution. Studies have shown today's OTT players and/or social media companies can introduce a new service in as fast as a day². While this may seem extreme, particularly given that operators will usually work with external parties (vendors) for development of new capabilities, the industry needs to transform to a much more rapid new function delivery and deployment model. HPE ContextNet provides its customers the tools needed to approach the agility observed by the OTT players.

- Resource allocation/reallocation as dictated by market: HPE ContextNet can scale one function and scale out another as the network traffic and business requirements dictate. For example, among many services that the service provider may have offered, when one or more seem to be more successful (e.g., higher revenues), then a Gi LAN function that is critical to that service can be expanded and another reduced, temporarily or permanently.
- Service providers use marketing forecasts to plan and dimension their network for several years. This stringency, which is not applicable to VMOs, can have financial consequences for their business. Errors in forecast in traffic volume (or customer acquisition) can lead to higher (or inadequate) initial investments resulting in idle resources or lost opportunities.

Scenario under Analysis

The scenario selected for analysis depicts the decision fork that operators with an existing PMO infrastructure are facing. The operator wants to introduce new services to a forecasted increase in its subscribers. Figure 4 illustrates this decision fork for expanding their Gi LAN segment.

² <http://www.cs.huji.ac.il/~feit/papers/FB13IC.pdf>

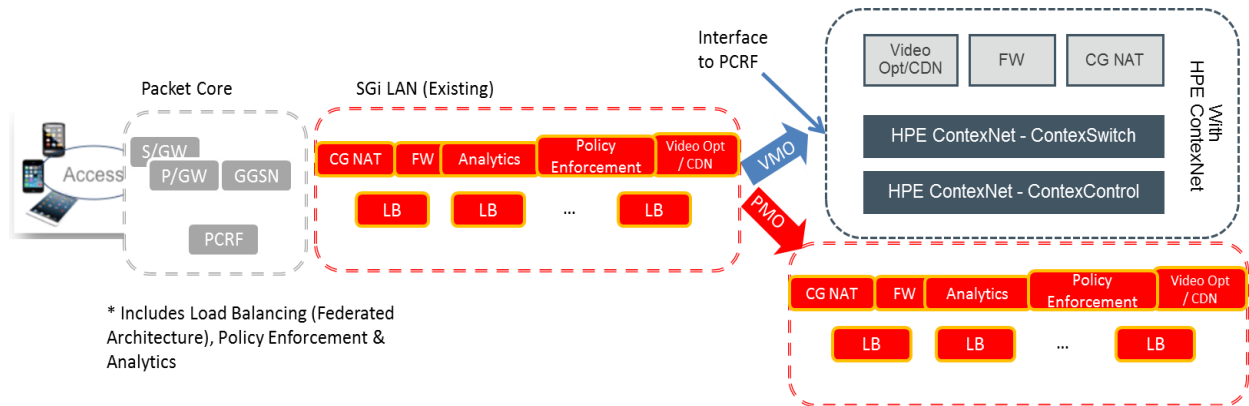


Figure 4. Network Diagram of Competing Solutions VMO versus PMO

The main premise of the scenario is expansion of the Gi LAN based on two choices: stay the course with expanding the existing infrastructure with a vendor that is offering a purpose-built appliance (PMO) or virtualize the expansion with a solution from a vendor that can be a replacement for the purpose-built appliances with virtualized appliances, i.e., VMO.

In this analysis, a PMO solution comprised of a set of purpose-built chassis and service cards (with 1x10GE I/O) to accommodate for Gi LAN services was selected. For simplicity, it was assumed that each card provides one Gi LAN service and operates at capacity. In some solutions, a card may provide more than one function, but this comes at the expense of capacity per function. In some implementations of PMO, different functions are deployed from independent vendors, encouraging best of breed, but in PMO this makes the network more complex. For this analysis, the PMO is based on an integrated solution (with one chassis/vendor capable of running multiple functions in an integrated manner), and the calculation is based on the assumption that a given function is deployed on a per card basis, with each card working at advertised capacity. The cards are assumed to be serially connected and traffic goes through all of the functions before they exit the Gi LAN (In alternate implementations the traffic may not go through all the cards, and more functions can be provided in software on one card, but then each card would operate at lower performance points). New load balancers must also be introduced to steer the traffic. In this mode, integration and rollout of new hardware and software is carried out with the existing Gi LAN components. The expansion in this case needs upfront planning and will be dimensioned based on the forecast given to the network designer and architects. It is designed for peak traffic whether the forecasts are on target. The paper has calculated the potential cost of deviation in a later section.

As in any network expansion, the service provider will go through a series of steps to integrate, test and verify the additional resources. It is assumed that the service provider is already adept in using the PMO and the integration/testing/verification and rollout take less time than the VMO case. This is true only for the first year of operation. After that, the integration and rollout for VMO declines relative to PMO.

HPE ContextNet solution (VMO) offers several immediate advantages relative to PMO:

1. Massive scaling for a high-performance solution, for example, 100 million subscribers, 1 million virtual network functions, 50 million PPS.
2. Pre-integration of three functions that are inherent in HPE ContexNet: load balancing via federated HPE ContexNet Controller, Analytics, Policy Enforcement, which contribute to reduction of network complexity.
3. Ability to dimension the network on demand without over/under provisioning (elasticity)
4. Ability to insert/remove functions dynamically without risk to existing services (thus reducing deployment cycle times and increasing agility)
5. The use of high-performance, lower cost COTS hardware to reduce capital expenditure (capex) and the ability to reduce operation expenditure (opex) with
 - a. Power consumption by calendar-shutdown of servers deemed unnecessary for a specific function or reduced traffic expectations
 - b. Network Operation Center costs per reduced labor and network automation.
 - c. Lower service maintenance cost per lower-cost servers.
6. Support for P+V packet treatment for overlay network to further lower network architecture complexity and monitor the entire Gi LAN operations.

In using the HPE ContexNet, the initial integration with the existing infrastructure will take longer than the PMO per hybridity of the infrastructure. In addition, Gi LAN functions will now have to be integrated further the assumption is that the same Gi LAN software that is used in the PMO solution will be virtualized and integrated in an industry-standard server. This initial investment in network rollout is not lost on anyone. The following years, any VMO integration will require lower time frames to complete relative to PMO. As previously mentioned, network rollout must be considered as part of our any capex calculation.

Capex Calculation Variables

The capex calculation for the PMO included hardware (chassis, blade, accessories where applicable, for example, interface ports), software licensing per Gi LAN function, and cost of network rollout. For the VMO, costs of industry-standard servers, Gi LAN VNF software, HPE ContexNet software modules (ContexControl, ContexSwitch, and ContexMap), and network rollout were considered.

OpEx Calculation Variables

Opex variables used were the same for both and consisted of:

- Engineering Facilities Installation
- Solution Life Cycle Management
- Solution/Product Support
- Network Operation Center
- Power Consumption and Real Estate Costs

Based on the forecast given by its market research team, the operator has to make a decision whether to:

1. Continue expansion based on purpose-built appliances (PMO) or

2. Expansion based on a virtualized mode of operation (VMO)

PMO System Configuration

The best integrated, purpose-built system available from a leading PMO vendor was selected for this comparison. It offers a chassis that can host a total of eight blades, each having 8x10GE I/O interfaces. Each blade is capable of running one single instance of a Gi LAN service, for example, CG-NAT. The PMO system requires that a set (one or more) of blades must be backed up by a backup blade for an N+1 redundancy configuration. In addition, a set of chassis must be protected by another one, again for an N+1 chassis redundancy configuration.

Other assumptions, for TCO calculation, on the PMO hardware and software configuration:

1. Assume maximum throughput for each blade (depending on the Gi LAN software).
2. Assume the most cost-optimal software/hardware combination that the vendor offers: blade that is bundled with Policy Enforcement may offer this cost optimization relative to purchasing them separately
3. The Network Roll Out (NRO) variable was also considered in the capex calculation. In general, this variable may have been ignored in industry, but it is important to consider. In general, the first year of NRO is to the advantage of PMO, because we assume that the labor for NRO is familiar with this step, having it done in the past.

VMO (HPE ContextNet) System Configuration

ContextControl, ContextSwitch, and ContextMap were used to configure this solution. Each module was assumed to be integrated in a separate HPE industry-standard server.

For the hardware, HPE DL380 servers were selected with these assumptions:

1. 1x 40GE ports with a maximum of 2 ports per server
2. Maximum of 20Gbps throughput per virtualized Gi LAN service, integrated into DL380
3. Two Virtual Machines per DL380

Gi LAN Services: Traffic Profile

The study assumed the following Gi LAN services and percentage of traffic that goes through each:

1. CG-NAT and Firewall – 100% for both PMO and VMO
2. Load Balancers, Policy Enforcement, Analytics – 100% for both however, for HPE ContextNet these are integrated within the various modules of ContextNet, therefore, no additional server hardware is necessary to host them
3. Video Optimization /CDN: 20% for both

TCO Results and Cash Flow Analysis

The Gi LAN network dimensioning was based on the following:

Scenario Assumptions: 5 Years	
Number of additional subscribers (expansion) in Year 1 with zero rate of growth for the next 4 years	4000,000
Peak bandwidth per subscriber in Year 1 ³	35Kbps
CAGR for peak bandwidth for Years 2–4	44.4%
Percentage of downlink traffic for Gi LAN	24%
Percentage of uplink traffic for Gi LAN	76%

Table 2. Assumptions for the Scenario under Analysis

The charts show the TCO advantages of the VMO solution, highlighting the lower TCO of the virtual solution for a small operator with 4M users. It needs to be pointed out that in addition to the savings VMO comes with level of agility that is very much needed and not possible with the PMO.

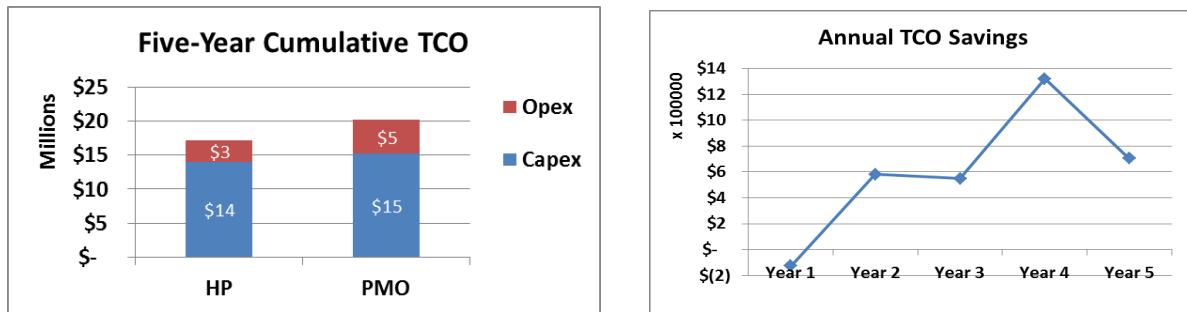


Figure 5. Five-Year Cumulative TCOs and Annual TCO Savings for HPE ContextNet Solution




Figure 6. Cumulative Cash Flow for Hewlett Packard Enterprise Solution

A total of a \$3 million is saved on TCO by selecting the VMO solution, giving an ROI of 114 percent over five years. Although TCO saving is important, the potential for new service creation and delivery combined with the ability to scale the number of services within a period of time by using the VMO can be deemed even more important than TCO savings.

³ [Forecast of Mobile Broadband Bandwidth Requirements](#), from ACG Research and data given by service providers

New Service Creation: Definitions and Analyses

A major incentive for an operator to virtualize its network is the promise of the ability to create new and existing services that are an order of magnitude higher in “service agility” and “service scale” relative to PMO. These two translate into faster time-to-market, which in turn map into capturing a higher market share and increase in revenue volume. Or alternately, for an operator with a PMO infrastructure, they translate into **opportunity costs**.

The two diagrams below depict the differentials between the two approaches for the steps that are taken for creating and deploying a new service(s). The steps are at a level of abstraction as each requires a series of tasks not shown but are estimated on a collective time frame basis.




Figure 7. Steps and Timeframes Required to Create a New Service with PMO




Figure 8. Steps and Timeframes Required to Create a New Service with VMO

As previously discussed, the “Relative Agility” between the two approaches is R.A. $(44/139) \sim 1/3$ or creating a new service via VMO can be accomplished in 32 percent of the time that it would take the PMO to create the same service. In nominal value, this is about four person months.

Task	Definition
1. Service Evaluation Groups to Assess Service Idea	Before any new service is considered for creation, the service provider (SP) creates a task force to evaluate the service and whether it fits the SP’s business Key Performance Indicators (KPIs). Because the full cycle for the PMO is historically long, the SP has to ensure a high probability of success for the new service. Cost of failure could be high. For the VMO, the cycle is much shorter and the cost of failure much lower; therefore, the SP with the VMO infrastructure can make this decision in a much shorter time-frame.
2. Resource Capacity Scaling, for example, Policy Controllers, Hardware	Assuming that a new service is accepted for creation, the SP now needs to assess and assign new (or an existing one if there is enough capacity left) chassis, blades, and ports. This is rather a cumbersome step for the PMO. For the VMO this is a straight-forward orchestration and provisioning step.
3. Billing Setup	This is an important step. It is most likely similar in the length of time it takes for both approaches.

4. Service Policy and Gi- LAN Configuration	Similar to Step 2, this is much quicker for VMO per ease of provisioning.
5. New Service Catalog Creation	Since this is a new service, the time it takes to create a catalog for it takes about the same time frame for both PMO and VMO, e.g., as new service profiles are created and inserted in a database; however, subsequent invocation of this service is much faster in VMO than PMO.
6. Service Creation	The final step is the actual service creation, which includes full integration, testing and verification of the new service into the system. For the VMO case, this can be accomplished much faster.

Table 3. Steps to New Service Creation and Readiness

Given the time frame differentials in the “new service creation,” we use this information to assess the revenue potentials for two new services between PMO and VMO and distinguish the economic differences.

Services under Analysis

The study considered two new services. They both demand a premium and an ARPU is assigned separately to each.

1. Offer streaming video/audio (V/A) services
2. Bundles of unlimited access for a small fee for specific Apps, such as Facebook access

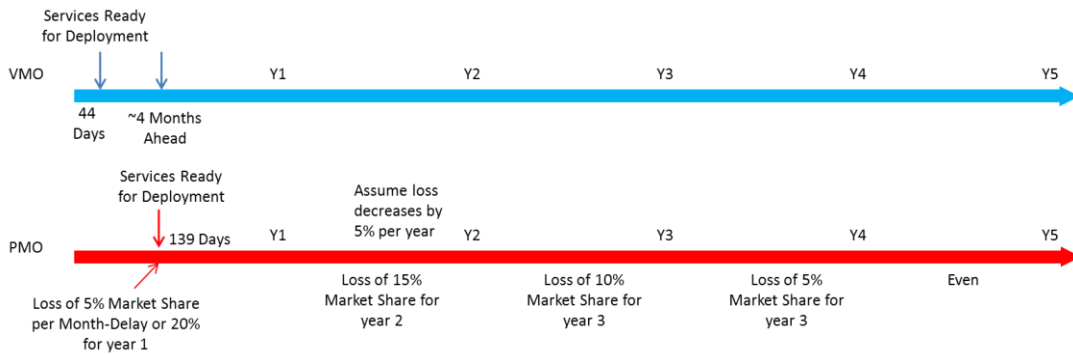
The assumptions for the new services are given in the Table 4.

Service Scenario Assumptions	
Number of subscribers signing up with V/A streaming in Year 1	500,000
Rate of growth in subscriber sign up for V/A next 4 years	5%
Number of subscribers signing up with Fee-based App streaming in Year 1	700,000
Rate of growth in subscriber sign up for next 4 years	8%
ARPU – video/audio streaming	\$10 per month per subscriber, no price increase, annually
ARPU – Fee-based App	\$5 per month per subscriber, no price increase, annually
PMO: Traffic flow through CGNAT, FW, Policy Enforcement, Analytics and Load balancer	100%

Table 4. Service Scenario Assumption and Assigned Values

The VMO infrastructure has a four-month advantage in TTM relative to PMO. This fact is used to make a few reasonable assumptions:

1. Because of slower TTM, we can assume that the PMO infrastructure loses an average of 5% of subscribers per month for each services that it is late relative to VMO. In this case, a 20% market share loss will be attributed to the PMO.
2. The loss in bullet #1 decreases by 5% after year one and per year thereafter as the service



offering matures and the subscribers begin to sign up regardless of the infrastructure.

Figure 9. HPE ContexNet Advantage in new Service Deployment Timeline

For VMO, a cumulative nominal advantage of about \$54.28 million is realized over five years, shown in figure 10.

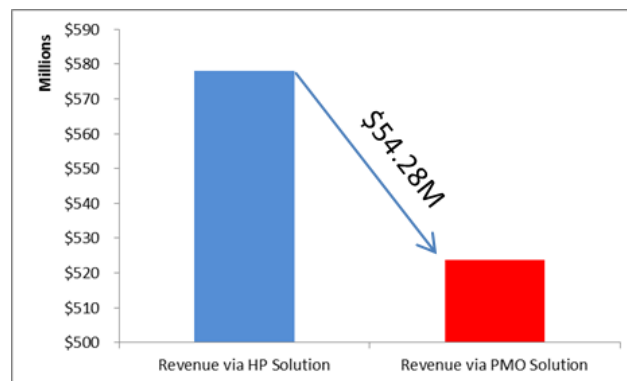


Figure 10. Cumulative Revenues per Solution

It is also important to consider the revenue generation levels relative to the TCO for each year. Figure 11 shows this for both VMO and PMO. The graph shows the relative Revenue to TCO ratio per year. It is normalized based on the “year 1” of HPE ContexNet. For example, if one dollar is the ratio for the HPE solution for the first year, PMO only generates \$0.82 for the same year. For year two, the Revenue to TCO ratio for HPE is 4.89. Similarly, for the PMO, this figure is 2.98. The chart is a representation of return on marginal efficiency of VMO.




Figure 11. Revenue to TCO Ratio per Solution

Other Revenue Analyses

Another major incentive for selecting a VMO based infrastructure is the scale with which service providers are capable to deploy services. The scale has changed the game for the service provider. Instead of a handful of services, say per year, they are able to deploy tens or hundreds of services if they choose to.

A sample analysis was done based on the following assumptions:

Service Creation Assumptions for a Five Year Period	
Average number of services offered per year for VMO	50
Average number of services offered per year for PMO	2
Probability that a service is successfully accepted by the market for HPE PMO	10%
Probability that a service is successfully accepted by the market for HPE VMO	100%
ARPU used per service	\$1 per month

Table 5. Service Creation Scale per Solution

A 100% probability of success was selected for the PMO solution since the operator needs to spend more time in vetting the solution at the onset, as the cost of failure could prove to be high.

Cumulative Revenue per 50,000 subscribers that sign up for these new services (we assume that each year the VMO operator brings 50*10% = 5 new successful services to the market, whereas the PMO operator only brings 2 new services:

Cumulative Revenue (C.R) over five years =

$$\sum_{n=1}^5 50,000 * \# \text{ of services deployed} * \text{probability of market acceptance for the service} * \text{ARPU} * 12 * n$$

Therefore, the cumulative revenue for each case is:

C.R VMO = \$45M per increase in SP's service portfolio

C.R PMO = \$18M

$(C.R VMO) - (C.R PMO) = \$27M$ for every 50,000 subscribers who sign up with a set of new services.

This, of course, is one sample calculation that shows the advantages of service scaling with VMO. There are two other major advantages (not included the above calculation) for the operator deploying the HPE ContextNet over the PMO operator:

1. It has the ability to tweak and improve a new service that may not have been a total failure but could become successful per subscribers' feedback. The operator has the ability to rapidly modify the service to suit the subscribers' needs and wants, and turn it into a successful service.
2. It has the ability to constantly improve a successful service based on subscribers' feedback which will result in a more satisfied subscriber base and lower the probability of churn. In addition, as subscribers experience better service, they are more inclined to sign up with other new services that the VMO operator offers.

Depending on the service provider's internal resources (labor and equipment) and the demographics that it serves, more services can be deployed and tested. If any of these services fail to attract the volume of subscribers that the service provider is targeting, it can be withdrawn quickly and with a much lower cost of failure for the VMO case.

Cost of Overprovisioning

Another incremental cost to consider for the PMO based infrastructure is the cost of over-dimensioning the network. This is not the case for VMO per its inherent elasticity, scale up/down, and scale in/out. Dimensioning the network for the PMO case is based on the forecast that is given to the network designers via market analysis and research data. If one of the important data points, for example, traffic volume, is off by only 3 percent, this will result in a waste of \$1.3 million over five years. This is another source of an opportunity cost for the operator with the PMO infrastructure, as the funds could be used for other purposes.

Conclusion

Network operators have been striving to generate new revenue streams and at the same time contain the cost of their network assets and operations. With the advent of Network Functions Virtualization (NFV), software defined networking (SDN), and service automation, operators expect to solve these challenges, alleviating cost pressures while benefiting from new innovative business models and services. NFV and its supporting technologies are innovative and disruptive architectures that promise to cut operators' capital and operational costs and enable faster time-to-market for services and, consequently, higher profit margins.

Although the TCO advantages of a VMO based infrastructure are naturally attractive, the variables that enable operators to increase revenue, benefit from clear competitive advantages, and create new



services with the agility that was not feasible with PMO infrastructure must be given more weight than cost savings.

ACG Research conducted an analysis of a typical scenario that a service provider is facing in migration to VMO based Gi LAN network segment as it tries to expand its network scale. The scenario assumes an existing Gi LAN infrastructure and the service provider has a choice to continue with PMO versus an HPE based VMO (HPE ContexNet) deployment. The study found a reduction of TCO level and a substantial gain in new revenue streams and volume based on only two new services over five years, as a result of faster TTM for a new service creation. The study also found that per increase in scaling in new services, additional cumulative revenue of \$27 million can be realized over a five-year period for every 50,000 subscribers that sign up with a service provider's increased service portfolio. In addition, for every 3 percent margin of error in traffic volume forecast, close to \$1.3 million was wasted per over-dimensioning the network infrastructure.

About ACG Research: ACG Research is an analyst and consulting company that focuses in the networking and telecom space. We offer comprehensive, high-quality, end-to-end business consulting and syndicated research services. Copyright © 2015 ACG Research. www.acgcc.com.