EXECUTIVE SUMMARY

Communication service providers (CSPs) are active participants in delivering cloud-based applications and flexible, scalable networking services based on virtualized network infrastructures. A critical success factor in supporting these services is the efficiency and speed with which CSPs can create and deliver them to their customers. A widely accepted belief among CSPs about achieving this efficiency is that their future infrastructures should be designed using architectures based on the abstracted, software-defined, general-purpose infrastructure philosophies pioneered by hyper-scale cloud and application delivery companies. These would provide the elastic and adaptable foundation their future service offerings require.

Although this strategy holds great promise, CSPs have important choices to make in how they pursue it. One of those choices is in the type of hardware they employ in enabling these diverse, software-driven offerings. There are significant differences in how alternative hardware platforms are designed and how they fare in enabling the efficiency, versatility and scale CSPs require. Options range from using openly architected, standards-based solutions built for integration and versatility in a software-driven world; to solutions built using significant amounts of proprietary design and which are more challenging to integrate in an open platform environment; and platforms developed using a bare essentials white-box approach.

In its use of a standards-based, open architecture philosophy for its CSP portfolio, HPE has had a keen eye on the advantages this approach can deliver for the CSP. Still they need to be articulated and quantified. To accomplish this, HPE engaged with ACG Research to compare the total cost of ownership (TCO) of the alternative approaches in CSPs' deployments. This report describes the results of that analysis, which demonstrate significant economic advantages for CSPs when using the open architecture, standards-based approach in their core and edge deployments for both NFV and cloud-based applications.

Key Findings

- Superior TCO + TTR in CSP service delivery is achievable using openly architected, standards-based physical system infrastructures in their core and edge deployments
- In an analysis of three alternative designs ACG found the cumulative 5 year TCO of open architecture designs to be 19% lower than the TCO of proprietary implementations
- Opex differences are more pronounced. Open architecture designs create 29% lower opex than proprietary ones and 86% lower opex than designs using off-the-shelf white-box platforms
- Time to results (TTR) of deploying open architecture designs is 58% faster than proprietary designs and 72% faster than bare bones white-box-based solutions
- The sources of these advantages are the simplicity, efficiency and flexibility of solutions using the open architecture designs

SUCCEEDING IN A CLOUD-NATIVE, OPEN ARCHITECTURE WORLD

Increasingly, application and network services are being delivered from a powerful infrastructure that blends both advanced communications *and* advanced computing technologies that simply deliver what end users need, when and where they need it. To accomplish this, both telecom and application providers are leveraging a streamlined infrastructure based on versatile, general-purpose computing using workload abstraction and elasticity techniques once pioneered in cloud computing installations and now being used for service delivery in virtualized networking, distributed and federated clouds.

A premium in this model is placed on achieving a *blend* of *lowering the cost* of delivering a given service while *maximizing the flexibility* with which a broad, evolving menu of services can be supplied. Clearly, large numbers of infrastructure elements must be deployed in a variety of sites, and yet the use of those elements must be optimized using efficient workload placement and management techniques.

Although a significant amount of this simplification can be derived from developing software applications that are flexible and adaptable, these goals have also created an incentive to design both open systems management data models and open communication protocols for managing the hardware of compute, storage and networking that in their own way are as extensible and scalable as the applications they are supporting. The goal is to create an efficient, elastic, and adaptable hardware infrastructure that can meet the performance and elasticity requirements of the networking and cloud-native applications running on top of it while maximizing the efficiency and return on investment (RoI) that service providers require.

For communications service providers (CSPs) playing a vital role in delivering these services, these objectives create the goal of developing an agile service delivery platform that stretches from core IT and network service delivery sites to the edge of their serving areas in smaller installations such as central offices, aggregation sites, mobile network radio towers, and in some cases, customers' premises. From the management and operational efficiency point of view, it creates a demand for consistently defined, open and standards-based elements that can be managed using well-structured, scalable automation. From the versatility point of view, it *also* creates a demand for a well-structured scalable management software design that can integrate new services efficiently.

This has produced a demand for automated, software-defined, and model-driven infrastructures. They are automated to achieve simplicity, efficiency and scale. They are software-defined to provide extensibility and versatility in the applications they deliver; they are model-driven to allow for consistency in the implementations that support them, which helps achieve both efficiency and agility at scale. Efficiency is achieved by not having to re-design the foundations of the solutions for each new functional requirement. Agility is achieved by being able to accommodate a variety of compositions and form factors in a deployment while not altering the fundamental mechanisms of configuration and control that the deployments as a whole require.

One *could* argue that this set of requirements is not entirely new. Although at a certain level they are not (for example, we have used SNMP and operating system command lines for several decades to work at accomplishing some of these goals), it is the scale, the versatility, and the demand for even greater efficiencies that have created the incentive to create designs that are more nimble and aligned with the web-scale techniques that have been used in cloud providers' deployments more recently and apply them to the management of CSPs' compute and virtualized networking infrastructures to achieve the efficiencies they require. The effect of these developments is to increase the scale and diversity of infrastructure that a given operations team can support while concurrently expanding the number of revenue earning services CSPs can deploy at a superior RoI.

In parallel with open-source initiatives such as OpenStack and standards-development efforts such as ETSI's NFV, these goals have played a part in motivating the development of infrastructure management information models such as the Open Data Protocol (OData[1]), useful in defining and managing physical system infrastructures such as servers in a consistent and extensible way; and Redfish and Swordfish scalable platform management specifications[2], which allow for those same infrastructures to be interfaced with, accessed and managed from a diverse array of management applications and devices by their operators. Indeed, to fulfill the requirements of scale and versatility that CSPs must meet moving forward, employing solutions that use these open, standards-based models will be a critical factor in their success. Employing them will allow CSPs to not only embrace the cloud-native, software-defined frameworks their applications require, but will also enable CSPs to achieve the level of efficiency and extensibility in their service delivery infrastructures, which will help them maintain their competitiveness moving forward. In short, leveraging model-driven solutions in their software-defined infrastructures will be a critical success factor in CSPs' future service delivery efforts.

HPE'S PORTFOLIO ALIGNED WITH THESE OBJECTIVES

HPE is a company that has recognized the importance of enabling CSPs' computing, storage and networking infrastructures with such standards-based, openly architected, model-driven capabilities and has moved proactively to contribute to their development. HPE is not the only company participating actively in these developments, but it has been noteworthy for the *extent* of its participation in them, as well as the substantive technical and leadership contributions HPE has made to the various communities' progress. This has been true in the Redfish, Swordfish, and NFV standards development initiatives, as well as in open-source community developments that have a direct effect on CSPs' ability to create more agile service delivery infrastructures (including the OpenStack, OPNFV, ONAP and CNCF/Kubernetes projects, to name a few of the more prominent ones).

A significant result of this proactive posture has been for HPE to incorporate key elements of standards-based and openly-defined functionality into its server, storage and networking portfolio focused on CSPs' use cases and deployments. Infusing its products with these capabilities helps those platforms meet CSPs' requirements for simplicity, efficiency and scale and at the same time provides the foundation for their agile service delivery infrastructures in NFV, cloud IT/application, and (increasingly) edge computing environments.

Clear examples of these designs exist in HPE's compute portfolio (including the Proliant, Edgeline, Cloudline and Apollo families) whose elements embed Redfish data modeling and API technology within them to support significant levels of automation, extensibility, and software-driven operation within their

[1] See: https://www.odata.org/.

[2] See: https://www.dmtf.org/sites/default/files/standards/documents/DSP0266_1.6.0.pdf for Redfish; and https://www.snia.org/sites/default/files/technical_work/Swordfish/Swordfish_v1.0.6_specification.pdf for Swordfish.

target use cases (such as NFVI for CSPs).[3] Additional examples exist in HPE's storage and networking product ranges with respect to simplified data modeling and management, but the underlying purposes and the resulting outcomes of simplification, versatility and efficiency in deployments are the same.

These implementations have not been undertaken in a coincidental or randomly occurring way. They have been done because the value of the infrastructure the CSP acquires when the infrastructure incorporates this functionality is higher than the value the CSP can achieve when its infrastructure is less consistently, openly or extensibly implemented. In other words, there is business value in the approach. Although appreciating that this value can be compelling, it becomes important for an adopter of the approach to quantify this value for the use cases being supported.

HPE DECIDED TO COMPARE THE ECONOMIC RESULTS ACHIEVABLE USING EACH OF THE TOP DESIGN ALTERNATIVES AVAILABLE TO CSPS IN THEIR DEPLOYMENTS

To obtain this data HPE decided it was imperative to analyze the use of its portfolio in a top representative use case it is addressing with CSPs as they evolve their NFV, cloud and edge computing initiatives. It also determined it would be useful to compare the results CSPs can obtain when using HPE's portfolio to the results a CSP can achieve when using either of the two leading alternative approaches to designing an NFV infrastructure that are available in the market currently and to summarize those outcomes. To give the analysis a measure of independent assessment, as well as to benefit from the experience of a third-party analyst in considering in the target use cases, HPE approached ACG Research to investigate the total cost of ownership (TCO) of the HPE solution compared with two representative and alternative designs in providing the NFVI for a prominent CSP use case (in this case, the virtualized evolved packet core of a mobile operator's network).

ACG Examined the Impact of Each Approach on CSP Infrastructures in a Representative Use Case (Deployment in a Tier 1 CSP's Virtualized Evolved Packet Core or vEPC)

ACG and HPE chose the vEPC because of its prominence among the use cases on which CSPs are focused in the deployment of their virtualized network services, and because it is fully representative of the types of requirements CSPs must address as they create a more agile and elastic infrastructure foundation for the delivery of enhanced services moving forward.

To give the analysis a realistic footing, we employed the example of a Tier 1 communication service provider's vEPC deployment as we can project it to evolve over the course of the coming five years.[4] As the technology that is currently available for larger scale deployments is based on 3GPP 4G or LTE, we

---

[3] See for example, https://www.hpe.com/us/en/servers/restful-api.html for a description of the Redfish API in HPE's Integrated Lights Out server management functionality.

[4] The dimensions in the analysis are based on a synthesis of ACG's work in service provider use cases over the past five years, including cases recently analyzed that factor in considerations similar to those we are analyzing in this model-driven infrastructure analysis. The identities of the CSPs whose environments are contributors to this representative profile, as well as the identity of the vEPC NFV solution supplier we make references to in our calculations, are kept anonymous at the request of those companies. Their dimensions, however, are applied consistently to each of the design alternatives we have analyzed in this project to ensure a level playing field of requirements to be addressed by each solution design.

evaluated the dimensions of an LTE based vEPC. We used the scale of the representative network as the baseline and incorporated the dimensions of a currently available VNF solution offering for the vEPC to present the same functional profile to be supported in each of the cases of the alternative infrastructures.

Moving forward, the introduction of 5G into operator networks will impact how their infrastructures are composed. There will be an evolution of the cores of the networks as 5G is introduced in either stand-alone (SA) or nonstand-alone modes. There will be an introduction of new site types, such as virtual RAN (or vRAN) nodes, in which underlying physical infrastructures of the type we are analyzing in this project will be introduced.

Many of the considerations we have examined in this analysis will be pertinent in CSPs' considerations as they prepare to support both LTE and 5G services. For this analysis, however, we maintained our focus on the offerings in CSPs' core network sites. The conclusions we have made in that context will have increased relevance in the larger and more distributed environments the CSPs will have in the future.

The dimensions of the LTE environment we considered in the case of offerings in each infrastructure deployment are highlighted in Table 1.

Attribute	Value in Year 1	Value in Year 5
Total # of Subscribers Supported in This vEPC	5 Million	5.6 Million
% of 3G Subscribers	20	18
% of 4G Subscribers	60	69
% of 5G Subscribers	0	10
Average Sustained Rate per Session per Type of Sub (4G Example)	373 Kbps	1,656 Kbps
Total Traffic (in Gbps) through This vEPC	1,867	8,252
Average # of Transactions per Second per Subscriber category (4G LTE Example)	49,920	64,474
Total Transactions per Second across All Subscriber Types for This vEPC	68,820	82,764

Table 1. Dimensions of LTE Environment

Of greatest impact on this analysis among the variables is the *growth in total traffic* in the networks the infrastructures will support over the course of the analysis period as subscribers' appetites for using mobile applications continue to grow. The effect of this growth on core network sites (such as the vEPC) is the growth in the capacity needed to support the demand. This translates into increases in capacity of every infrastructure element involved, from the physical infrastructures on which we are focused in this analysis to the virtual system infrastructures and the VNFs themselves that provide the operators' networking service functions. Clearly, the efficiency with which this growth can be achieved has an impact on the economic outcomes CSPs can expect to achieve. Thus, this high-growth environment of expanding traffic provides an ideal case for analyzing the efficiency of the alternative designs we are comparing in accomplishing the operators' goals.

SUMMARY OF THE IMPLEMENTATIONS

To make our comparison we used a generally equivalent set of configurations representing the three deployment models that were chosen as the focus of our analysis. These three models are:

- the open, standards-based architectural model employed by HPE (which we abbreviate as OSBI for open, standards-based infrastructure in the remainder of the report);
- the closed, proprietary architecture employed by other physical infrastructure solution suppliers in their offerings (which we abbreviate as CPSI for closed, proprietary system infrastructure); and
- the unmodified, off-the-shelf white-box implementation employed by additional suppliers in their offerings (which we abbreviate as OWBI for off-the-shelf white-box implementation).[5]

We arrived at configuration sizes needed in each alternative by determining the amount of each resource required to support the VNFs and the related virtualized infrastructure management (VIM) and NFV infrastructure (NFVI) software required for the vEPC. This is based on the attributes of the VNF, VIM and NFVI software that have been implemented by the various suppliers with whom we have worked, and by applying the dimensions to the different infrastructure designs using the attributes of those designs, and of the representative Tier 1 CSP deployment we are referencing.

The configurations we analyzed include the compute (server), storage, and network (switching) components needed to support the vEPC. We based the characteristics of each element used in the configurations on the elements used in HPE's core NFV site blueprints.

HPE Telco Blueprints include server, storage and network elements designed to support a core NFV installation, along with the required rack and accessories needed for installation. In a vEPC like the one we have analyzed, deployment includes servers for management and control, compute workloads (such as for VNFs), and storage for both virtual system components (such as VM images) as well as operational data for the deployment (needed by VNFs and management applications).

The requirements for each category are distinct. As such the blueprints specify distinct central processing units (CPU), such as Intel Xeon Scalable Silver 4116 for management and control server nodes versus Xeon Scalable Gold 6152 for compute server nodes; distinct memory and storage arrays (varying the type and amount of HDD and SSD for the requirement); and switching elements to provide connectivity within and between the racks, as well as other network domains.

HPE products from the ProLiant Gen 10 server family and the FlexFabric switching family that we employed are shown in Table 2. They are included in HPE's core NFVI blueprint.

Type of Node in the vEPC	HPE Product Employed
NFVI Control and Management Nodes	HPE DL360 Gen 10 w Xeon Silver 4116 2.1 GHz 12-core
NFVI Compute Nodes (VNFs)	HPE DL360 Gen 10 w Xeon Gold 6152 2.1 GHz 22-core
NFVI Storage Nodes (Images and Data)	HPE DL380 Gen 10 w Xeon Silver 4116 2.1 GHz 12-core
NFVI Network (Top of Rack) Switches	HPE FlexFabric 5950 32QSFP28 Top of Rack Switch

Table 2. Representative HPE Products Used in the vEPC Configuration

[5] We keep the identities of the suppliers using the approaches that are different from HPE's approach anonymous in our report as our focus is on the economic implications of using the alternative approaches when deploying into the representative environment versus comparing the suppliers' offerings directly.

We used generally equivalent elements from the alternative suppliers to define the configurations in their cases. This means we selected server, storage and networking products employing as close to the same CPU, memory, storage and networking functions as is possible to make the comparisons. Clearly, not every feature and function is supplied in exactly the same way. If that were the case, there would be no real comparison to make. However, on the primary attributes involved, the configurations are aligned.

To illustrate the scale of the vEPC under analysis, the number of servers and switches we included overall is shown in Table 3.

Type of Element	# in Year 1	# in Year 3	# in Year 5
Total # Servers (All Types)	145	284	539
Total # ToR Network Switches	10	20	34

Table 3. Total Number of Infrastructure Elements in the vEPC

Metrics Employed in the Comparison

We considered the impacts of the approaches relative to operators' fundamental economic goals using both TCO and time to results (TTR) measurements. Combined, these give an indication of which alternative an operator is considering is likely to provide greater returns faster.

As an overall metric, TCO comprises both *capital* and *operating* expenses (capex and opex) involved in deploying a solution. In our scenario, the capital expenses include the costs of the servers, storage, network switches, and the rack-mounting infrastructures in which they are deployed. In tandem, the operating expenses for the alternatives include personnel costs in designing, testing, installing, and operating the infrastructures for the NFVI. Opex also includes components such as electrical power and floor space consumed, though these are, by and large, less impactful to the TCO than the cost of staff's time for doing the work. The *cumulative* nature of the TCO provides an integrated view of capex and opex over the course of a given length of time. In our case we considered the alternatives over five years.

TTR is a complementary metric to TCO. It identifies the length of time (in calendar terms) it takes to accomplish the objectives of a deployment. If one alternative enables revenues or savings to be achieved at an earlier date than another (because of the efficiencies it enables) its TTR is more appealing.

SUMMARY OF RESULTS

Here we present a summary of the results according to each category of cost and the overall effect each has for the CSP. As we will see, it is ultimately a combination of the efficiencies in opex and the resulting advantages in cumulative TCO that make the use of the open, standards-based designs the most advantageous of the three alternatives for the CSP.

First, consider the five-year cumulative TCO of the alternatives. In our analysis we determined the cumulative TCO of the standards-based, model-driven HPE NFVI to be 19% lower than that of the TCO of the closed, proprietary NFVI, and 7% lower than the off-the-shelf white-box-based implementation.

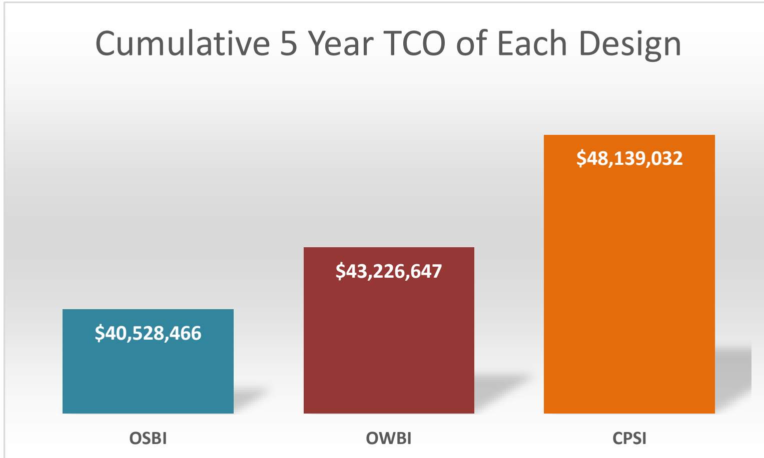

Figure 1. Cumulative TCO of the Alternatives

By definition, this represents the combined capital *and* operating expenses associated with each approach. Applying this outcome across a significant number of CSPs' sites, the absolute value of the savings would increase, to say, $80 million over five years in a deployment involving 10 sites of the same configuration size. One could view this amount as funds made available to other operating goals by choosing one approach versus the other.

As we dig further, however, we see that the agility and the efficiency advantages of the OSBI in opex become the more compelling consideration overall. This is based on an alignment of the characteristics of that alternative with CSPs' strategic requirements of supporting diverse new services and workloads flexibly and quickly at the lowest possible cost.

Let's unpack this according to the major components of the TCO. First, in Figure 2, as we might expect, we see the pure capex of acquiring an off-the-shelf white-box offering is lower than the capex of either the open, standards-based implementation or the closed, proprietary design. These differences are attributable to the design approaches used in each supplier's case. White-box suppliers concentrate on incorporating just the essentials needed to get a platform up and running in its elemental form (as an x86 server, a storage array or a networking switch). It is up to an ecosystem partner, an integrator or an end customer (such as a CSP) to supply more refined capabilities that will make the white-box ready for operation and deployment in its targeted case. By comparison, an OSBI supplier such as HPE has taken the steps to enhance the components of a configuration using standards-based, openly architected additions that, in the end, simplify and scale a deployment in the customer's environment, balancing both efficiency *and* choice for the customer. Finally, suppliers in the CPSI category while striving for enhancements and efficiencies as well, do so using fewer of the openly-architected, standards-based implementation options than the OSBI or the OWBI suppliers generally use. Hence, we see the differences in capex that we have identified.

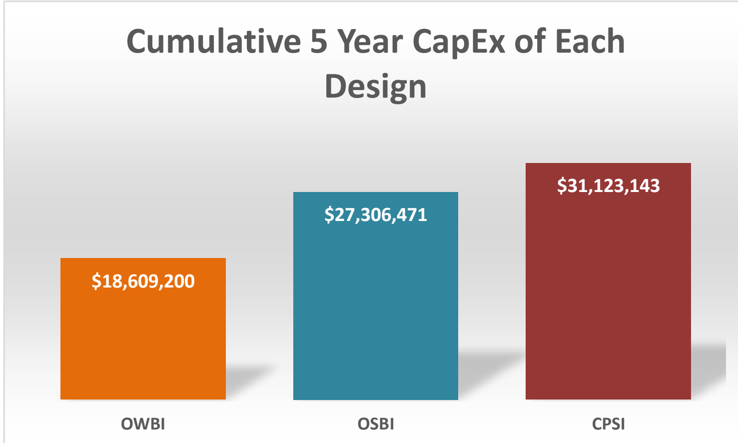

Figure 2. Cumulative Five-Year Capex of the Alternative Designs

It is when we expand the analysis to include the opex effects of the alternatives that we see the greater impact of the OSBI emerge. Figure 3 shows the cumulative five-year operation expense of the alternatives in the CSP's vEPC.

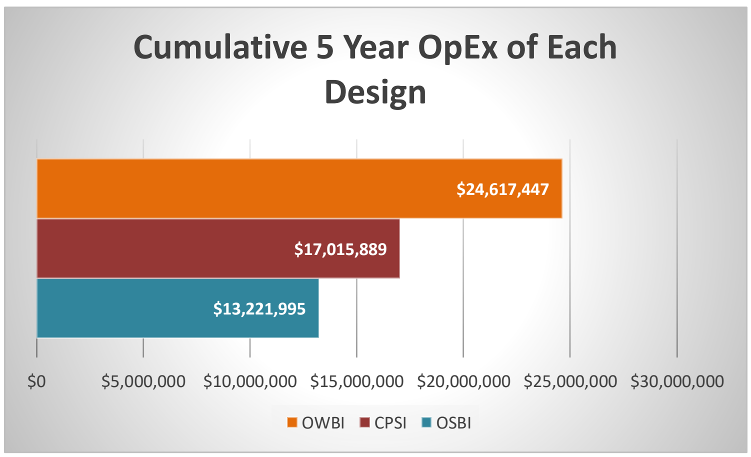

Figure 3. Cumulative Five-Year Opex of the Alternative Designs

Clearly, the OSBI is more efficient (cost effective) in deployment and operations terms than either the CPSI or the OWBI. Over five years in this use case, the OSBI is 29% more efficient than the CPSI and 86% more efficient than the OWBI.

Why is this the case, and where do the differences come from? The differences are based in the details of operational cost and how it is incurred. These are sometimes fuzzy to identify when comparing alternatives in the abstract. They come more concretely into focus when identifying how staff's time is spent when preparing to deploy a solution and when running the solutions in production. It is the specifics of the individual tasks and the aggregate of those allocations that comprise the differences in opex between the alternatives.

The sources of the opex differences we identified across the complete life cycle of deploying the vEPC are broken out in Figure 4.

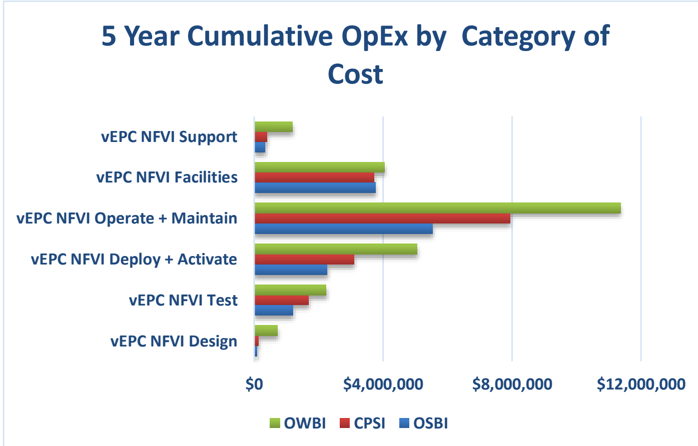

Figure 4. Cumulative Opex of Each Alternative by Category of Cost

As we can see, the differences are greatest in the areas of solution test, integration, deployment, activation, and ongoing operations and maintenance. This is largely because of the differences in support for necessary functions that are built into the OSBI solution at the start versus the amount of work to enhance, integrate and provide customized procedures in operations for both the CPSI and the OWBI solutions, in differing degrees, to make them each deployable, manageable, and extensible for the use case. These differences have an effect on the overall cost of deployment in each case (shown in both the opex and the cumulative TCO charts) as well as in their relative TTR.

Differences in TTR between the Alternatives

In parallel with the differing levels of TCO we calculated, we determined it is faster to make NFV environments ready for delivering revenue-generating services using the open, standards-based infrastructure design than using either the closed, proprietary infrastructure or the design based on bare bones white-box infrastructure components. This is because of the reduction in time enabled by the simplifications achievable using the standards-based approach, in conjunction with the guidelines of the pre-validated blue prints made available with it versus the comparatively more complex and time-consuming tasks required to make the proprietary and the bare bones white-box infrastructures ready for production use. Figure 5 shows the open, standards-based infrastructure enabling a TTR that is 58% faster than the deployment based on the design of the proprietary system and 72% faster than the deployment based on the bare bones white-box design.

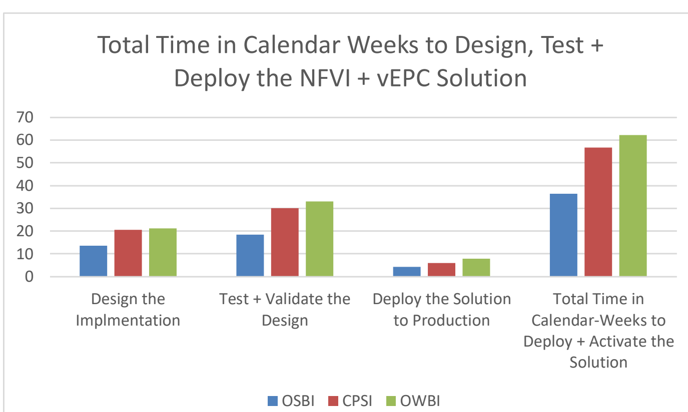

Figure 5. Time in Calendar Weeks Required to Deploy Each Alternative

The time required in calendar weeks to deploy a solution is a different metric than the total cost of the personnel time required to accomplish the deployment. Calendar time is consumed without specific reference to the number of personnel investing their time during that calendar period. Theoretically, there could be 10 person-weeks in personnel time invested to accomplish a given deployment phase in one solution's case, while the same phase of deployment may require 20 person-weeks of invested personnel time using an alternative solution. Those aggregations of time (which equate to the personnel cost in dollars to the operator of performing that phase) do not directly define the length of calendar time required to complete that phase in either case. The first instance might require 4 calendar weeks to complete; the second might require 6. Thus, the first instance becomes ready using 67% of the time that the second instance, but using 50% of the personnel dollars when doing that. The two metrics (opex costs referring to personnel time invested in accomplishing the work, and time to results looking at calendar time required to accomplish a deployment) are somewhat related to each other overall but are not the same in their calculations, and in what they impact. The TTR, in the end, determines how quickly the operator can begin receiving the benefits (either savings in operations, or revenues being generated by the deployment – or both) made possible by the implementation that is ready in the shorter length of time.

Diagnosing the Results

The differences in TCO and TTR a CSP can expect when deploying an NFVI based on each of these solutions are present because of differences in the design of each of the alternatives. We highlight the essentials of those differences in this section and describe their impact on the results.

The differences are traceable to the following attributes of their design:

- Use of industry standard data models and APIs internal to the operation of the products.
- Mechanisms for integration with external systems using industry standard protocols and data models.
- Extensive integration of these standards and models throughout the portfolio.
- Active participation in the standards bodies and open-source communities, accelerating the delivery of open solutions for NFVI and cloud-native application environments.
- Extensive collaboration with a broad ecosystem of partners to enhance usability of the portfolio in a range of CSPs' environments.

Additional detail on each will help clarify their impact on the results.

Use of industry standard data models and APIs in product design. This concept is susceptible to a variety of interpretations, depending on which part of the infrastructure one is referencing. In this comparison, we are focused on the efficiency of deployment, integration and operation of the underlying physical infrastructure of the NFVI. We are focused on the consistency and openness of the functions the platforms use with respect to installation, integration and management of their hardware resources across the life cycle of the deployment of the NFVI. The degree of openness and standardized representation that equipment uses to describe its configuration, communicate about its state, and provide information about its capacities and capabilities to the various layers of software that have an interest in that information as the operation of the deployment progresses is in what we are interested. Examples of such higher layer constituents of the underlying hardware include the VIM that is interested in placing workloads in the proper underlying resource and the monitoring software being used to understand the utilization and state of resources such as CPUs, memories and NICs within the NFVI.

A primary standardized data model and interface specification the computing industry has developed to facilitate the open and efficient supply of information with higher level functions about underlying characteristics of a hardware element is the Redfish Scalable Platforms Management Specification developed by the DMTF.[6] The Redfish specification and data model define a powerful and efficient mechanism for simplifying, streamlining, and automating the operation of underlying computing resources in an NFVI environment. Incorporation of the Redfish model exists in a variety of solution offerings CSPs can employ. The differences in the cases we have examined in this analysis are the degree to which the model has been employed, the simplicity and directness with which its capabilities have been made available to the managers of the related elements, and to the higher level software functions that also stand to benefit from the availability of the functionality.

HPE's integration of the Redfish specifications into its compute portfolio servers and its iLO resource management functionality provide a more extensive implementation of this functionality than is available in either of the alternative offerings. The contribution of its operations in initial installation, startup and discovery of configurations, supply of information in testing and operational management functions, such as performance and problem isolation/resolution, is material in creating efficiencies for the use of the HPE solutions within the NFVI. The integration of these functions consistently and elegantly into higher level management capabilities generates a faster time to results ranging from 20 to 50% across many of the tasks involved in validating the proper operation of the NFVI in the vEPC. This is made possible by the disciplined implementation of the standardized data elements and a standardized mode of accessing them by both human and software-driven managers that need to consume the information.

Use of industry standard protocols and data models in mechanisms supporting integration of the hardware infrastructure with software that needs to interact dynamically with the hardware when performing its operations. This is largely a portfolio completeness and open platform architectures point. The way individual products present their operation to other software entities has a significant effect on the speed with which the other software elements can be enhanced to support the hardware as it evolves and also affects the efficiency and scale at which the software and the hardware can be used with each other in important deployment environments such as NFVI and CSP application infrastructures. The more open, widely adopted, and consistently implemented in the integration interfaces are, the easier and faster it is to deploy, run, and get value out of the overall solution.

The best example of this is the use of the Representational State Transfer (REST) software architecture and protocol for interoperability of HPE's DL series servers with external applications along with the JSON data format for exchanging the data with the other systems. REST and JSON are significantly easier to integrate with and more widely adopted within the industry than, say, solutions using XML instead of JSON. The benefit of this emerges in the integration of the server infrastructure with higher level auto-provisioning, analytics and platform management software that needs to work with NFV deployments at scale. This capability provides a significant efficiency and scalability advantage for HPE's portfolio

[6] See for example, *Redfish Scalable Platforms Management API Specification, Version 1.6.0, 08-23-18*, developed and published by the DMTF: https://www.dmtf.org/sites/default/files/standards/documents/DSP0266_1.6.0.pdf.

compared to the alternatives we examined, which is especially meaningful in the size and complexity of the vEPC environments we are considering.

Extensive integration of these standards and models throughout a versatile NFVI portfolio. The virtues of designing standards-based functionality and openly-architected functions into an individual product (such as a given type and size of server or switch) are clear. What is even more impactful is designing consistently architected versions of those capabilities into complete ranges of complementary products that support a similar capability (such a running a VNF workload) and yet make it possible to support that type of workload efficiently, consistently, and at scale across a wide range of form factors, numbers of components, and configuration mixes required to be fit for the CSP's diverse installation requirements. Building a portfolio that is suitable for use across this full range of use cases and sites delivers an amplified level of benefit to CSPs as the operations needed to support that range of sites evolve. The fact that HPE has designed its support of functionality such as Redfish, REST, Netconf and other open platform capabilities into its DL series servers (more typically found in a core or virtualized central office), its Edgeline (EL) series servers (more typically found in a network edge location), its FlexFabric switches, and its related management software solutions in a manner that supports operation in CSPs' deployments at the scale we are referencing in this document is an indication that HPE has taken the value of this consistent implementation of standards-based and open architectural designs seriously and factored it into the development of the portfolio it is applying to CSPs' NFV and cloud application requirements.

Active participation in the standards bodies and open source communities accelerating the delivery of open solutions into NFVI and cloud-native application delivery environments. Any design team can decide to incorporate a given specification or open-source component into its solution that it believes will be useful to its users. That said, it is different when the development team for a portfolio is actively engaged in and leading the architectural specification and/or open-source development processes that have a bearing on the use cases for which that portfolio will ultimately be deployed. In HPE's case, its leadership efforts and ongoing participation in both the DMTF (related to the Redfish specifications), ETSI's NFV specifications groups, and emerging initiatives in ETSI and in the Linux Foundation related to the distribution of NFVI and cloud-native functionalities more broadly into CSPs' environments warrant being recognized as providing an advantage to its development efforts and its customers for a broad and comprehensive embrace of the frameworks that are created within those communities. Acknowledging this does not provide a guarantee of high-quality implementation or of uniquely useable solutions; that outcome only comes from executing the implementation of the deliverable solutions that embody the designs. However, at a broad level, HPE's efforts show evidence of not only providing leadership to the innovation and specification efforts, but also of incorporating the frameworks into its portfolio in support of CSPs' NFVI, cloud application and emerging network edge deployments.

There is participation by the suppliers of the alternative solutions we are analyzing in this report in certain areas of the efforts we have mentioned but by no means all and also not as extensively as is the case with HPE. In this sense, to the extent that HPE's contributions in these communities help bring capabilities into platforms using, say, Redfish or JSON or enhancements to NFVI infrastructures in how virtual system workloads are supported in an NFVI, the efficiencies we have noted in the HPE solution offerings related to the NFVI environment we are examining are the direct beneficiaries of those community contributions.

Extensive collaboration with a broad ecosystem of hardware and software partners to enhance the usability of the portfolio in a wide range of CSPs' environments. Aside from concentrating on its own portfolio of products to be used in an NFVI, the efficiency and flexibility with which HPE's portfolio can be used by CSPs is substantially enhanced by HPE's collaboration with a broad set of partners to enhance the ease with which its products can be integrated into a CSP's NFVI, such as the vEPC we are examining in this case. The value of this collaboration is nearly as impactful to the results CSPs can achieve as HPE's participation in the standardization and open-source communities.

For example, via close collaboration with a series of VIM and virtual systems infrastructure suppliers that are integral to the deployment of an NFVI, HPE's physical infrastructure elements can be brought to bear on a CSP's deployment in a variety of virtual software infrastructure implementations. This does nothing other than increase the flexibility the CSP has in considering how it should proceed. In the case we examined in this analysis, the VIM and virtual software infrastructure we used as a reference implementation for supporting the vEPC are supplied by one of the software suppliers to which we are referring. That close integration has an impact on the speed with which the NFV solution can be designed, integrated, tested, deployed and run.

A similar circumstance exists with suppliers of the VNFs that are needed for the use case a CSP is pursuing. In the case of the vEPC, HPE's collaboration with several suppliers of vEPC solutions provides both flexibility of choice and efficiency in deployments for the CSP. Again, in this case, the reference implementation we employed is supplied by a vEPC partner with which HPE has worked closely.

Finally, as is well known, HPE works closely with CPU, memory and other component suppliers in developing its server, storage and networking offerings for use in NFVIs. Especially in the case of its CPU and memory relationships and because of the deployment customizations as to core allocations and I/O optimizations that are important within the vEPC, HPE's close working relationships with both underlying component designers and with the OS and virtual infrastructure software suppliers we referenced, HPE is in a position to work closely with CSPs to tune quite precisely its configurations and deployments for the needs of the NFVI.

In all of these instances, the results of these extensive collaborations are to supply a blueprint and design for NFVI deployments that are comparatively more straightforward to validate and run and that provide substantial operating efficiencies based on the alignment of the components with their tasks. In the case of the bare bones white-box solution offerings, by comparison, there is not the depth of tuning and validation for the task as is the case in the HPE portfolio. And in the case of the more tightly coupled and more proprietary design, while there are ecosystem partnerships that are at work for that offering, the extent of those engagements and the efficiency with which implementations can be realized when bringing the solutions together in a multi-layered use case such as the vEPC is comparatively constrained relative to the breadth of implementations included in the HPE solution suite.

Advantages to CSPs in Using Openly Architected, Standards-Based, Software-Driven Infrastructures

We can see from the analysis there are clear advantages to CSPs for using open-architected, standards-based, and software-driven infrastructures in supporting the virtualization of their network infrastructures, as well as establishing an agile foundation for their delivery of cloud-native, widely available applications moving forward. Let us capture what those key advantages are and highlight how they benefit CSPs in these pursuits.

Advantages accrue to CSPs from the use of infrastructures embodying the characteristics we have identified on each of the following dimensions:

- Speed
- Efficiency
- Agility
- Flexibility
- Scale
- Competitive Advantage

Speed. CSPs can execute their operational tasks more quickly (as we have seen) using open architected platforms (shortening their time to results), and they can create and integrate new service and application offerings more rapidly (accelerating their time to revenue). This is based on the flexibility provided by an open, standards-based, and software-driven architecture, which allows for interfaces between modules to be implemented using well-known, consistently defined methods that can be adapted to the needs of each new service offering. This creates a wide range of options for CSPs to employ without imposing unnecessary constraints as to timing because the various implementations they are considering are each pursued using a proprietary design. In the open architecture plan, the barriers of time to achieving intended results are minimized.

Efficiency. Efficiency in speed of execution is obvious, but what we note here is, based on all of the attributes of the open architecture approach, an economic efficiency that shows up as *superior financial returns* from pursuing this approach compared with alternatives. This is, ultimately, the most impactful reward from adopting a given path. We note these returns in the form of cumulative TCO in this case. That superior TCO contributes to the CSP achieving an improved operating return in the periods in which the approach is being used.

Agility. By agility we refer to the CSP’s ability to address a wide range of service delivery opportunities confidently based on the diverse capabilities of an underlying platform that uses the open architecture approach. By being able to address the different requirements of a mix of applications efficiently using the software-driven, open architected attributes of an NFVI, which is implemented using these principles, the CSP can supply the widest range of customer-facing services on a foundation introducing the lowest overall cost of entry. With the growth CSPs are experiencing in offerings, ranging from streaming video to multiple categories of IoT and the distribution of intelligence throughout their networks to support them, deploying an infrastructure that has these agile properties is a critical requirement moving forward.

Agility also comes into play at the hardware level when the platforms are designed to accommodate it. For example, with the flexibility to introduce FPGA based hardware optimizations for functions such as transcoding and encryption or graphics processing units in parallel with CPUs, a hardware portfolio supplier delivers significant additional agility to its customers by being able to make those choices when and where they are appropriate while still incorporating them into a standards-based operational model.

Flexibility. Although the term may sound approximately the same as the efficiency and agility, what we mean by flexibility in this context is the opportunity that an open architected, standards-based infrastructure provides the CSP for choosing the type of VNF, application or management software implementation that best suits its needs for the implementation of its services. In real terms, it is possible to consider a wide mix of solutions, sourced in a variety of ways for meeting the requirements while maintaining the same underlying multi-tenant foundation in place to support each one. As time progresses, if it is crucial to consider changing to a new alternative in the case of a given application, that foundation can continue to be used (and evolved) while the applications using it are swapped out or evolved on their own.

Scale. In the CSP's environment, and especially those of mobile network providers, the ability for a chosen infrastructure to scale economically *and* reliably is a critical success factor in both 4G/LTE and 5G services. Scaling will occur in both core network sites, such as those evaluated in this study and in distributed edge sites, including virtualized central offices and in vRAN resource pools installed at various points in the operator's aggregation network. What has begun as a deployment of half-dozen racks in a core network site is quite likely to scale to three or four-dozen racks in those locations along with deployments into hundreds of remote installations across the footprint of the operator's infrastructures. In that future, having a standards-based, open architected and consistently implemented fleet of resources to manage to support the mix of services being offered will be of strategic value to the CSP.

Competitive Advantage. In the future, operators will compete largely based on how quickly and flexibly they are able to bring specific service offerings to customers within the deployment environments in which they are working, as well as in federation with ecosystem partners supplying their applications in close cooperation with the CSP. This means the faster new locations can be supported and the faster new applications can be integrated into a software-driven, open architected infrastructure, the more attractive a CSP will be to its end customers *and* its partners. Deciding to employ an infrastructure foundation that enables that speed, flexibility and efficiency to be achieved will be one key part of establishing a competitive advantage in delivering a given set of services for the CSP.

Openly Architected, Standards-Based and Software-Controllable Designs Enable the Cloud-Native Future CSPs Require

As we have seen in our analysis of design alternatives for use in the vEPC use case, there are clear and distinct advantages to CSPs from using an NFV and cloud application infrastructure that is built from the ground up utilizing open architected and standards-based mechanisms for instantiating and evolving their deployments. It is unmistakable that this architectural perspective yields advantages in both efficiency in life-cycle management of the installations, as well as superior agility and versatility in supporting the diverse networking and application workloads that are the essence of CSPs' revenue streams.

Moving forward the importance of using this open and agile architectural framework in supporting new applications in both edge computing and multi-cloud operational scenarios will only increase. In those scenarios the ability to support a diverse mix of application requirements dynamically and a large and highly distributed footprint and the ability to automate those deployments in a heterogeneous combination of service providers' domains will be competitive necessities for CSPs to embrace as they bring a larger, more diverse, and more valuable mix of applications to bear on their customers' needs.

The architectural and ecosystem development approaches HPE has embraced in the evolution of its portfolio for CSPs is firmly and directly aligned with these objectives. They are a strong example of how an openly architected and standards-based approach to providing an agile service delivery platform for CSPs in delivering their virtualized networking and their cloud-native application services provides superior value and return to the CSPs and their partners as they broaden their participation in the fully connected, cloud-native application delivery environment in which their customers are working.

Table 4 compares the cumulative capex, opex and TCO incurred by CSPs over five years when they deploy an NFVI for a virtualized EPC based on the OSBI, the CPSI, and the OWBI that we analyzed in this report.

Metric	OSBI	CPSI	% OSBI Advantage	OWBI	% OSBI Advantage
Five-Year Capex	**$27,306,471**	**$31,123,423**	**-13.98**	**$18,609,200**	**+31.85**
Five-Year Opex	**$13,221,995**	**$17,015,889**	**-28.69**	**$24,617,447**	**-86.19**
Five-Year Cumulative TCO	**$40,528,466**	**$48,139,032**	**-18.78**	**$43,226,647**	**-6.66**

Table 4. Comparison of Five-Year Capex, Opex and Cumulative TCO of the 3 Alternative Designs

Hewlett Packard Enterprise

About Hewlett Packard Enterprise
Hewlett Packard Enterprise advances the way people live and work
Learn more at hpe.com/dsp/infrastructure

intel

About Intel
Intel Inside®. Network Transformation Outside
The era of 5G-powered experiences starts today with Intel® technologies
Learn more at www.intel.com/networktransformation

ACG RESEARCH

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. www.acgcc.com.