



White Paper

Analyzing the Economic Value of Partner-Delivered HPE Converged Architecture 700 Infrastructure

By Adam DeMattia, Research Analyst

January 2016

This ESG White Paper was commissioned by HPE and is distributed under license from ESG.

Contents

Introduction	3
Executive Summary	3
Market Overview	3
HPE Converged Architecture 700: Economic Value Analysis	6
Methodology	6
Economic Value Model Overview.....	6
Cost Categories.....	7
Benefit Categories	7
Default Scenario	7
Summary of Results.....	8
Benefits Analysis.....	10
TCO Analysis	13
The Bigger Truth	15
Appendix A.....	16
Appendix B.....	17

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.

Introduction

Executive Summary

ESG was engaged by HPE to develop a detailed economic analysis of its partner-delivered HPE Converged Architecture 700 infrastructure platform. The analysis is designed to help IT organizations determine the fully-burdened costs and benefits of leveraging this pre-validated reference architecture compared with a “present mode of operation” (PMO) that reflects traditional component-based infrastructures integrated to support virtual environments. This analysis builds upon ESG’s evaluation of Converged Architecture 700, in-depth interviews with technical stakeholders at HPE, ESG qualitative and quantitative market research with IT decision makers, and ESG’s general familiarity with the adoption drivers and perceived advantages of converged and reference architecture platforms. This analysis is designed to provide prospective customers with a comprehensive picture of the potential direct and indirect cost and benefit drivers they should consider when evaluating an HPE Converged Architecture 700 investment.

As discussed in the following pages, HPE Converged Architecture 700 offers the opportunity for organizations to increase their IT and user productivity, while significantly improving application time to value and achieving TCO parity compared with traditional infrastructure approaches. **In fact, ESG’s analysis of a typical use case for HPE Converged Architecture 700 results in an impressive 156% ROI and a brief ~9-month payback period.** Traditional approaches to infrastructure virtualization that ESG modeled resulted in significantly lower ROI as a result of significantly lower expected benefit. For organizations struggling to overcome IT operational challenges tied to infrastructure complexity, embracing the automation and intelligence offered by modern integrated platforms can prove invaluable. Computing solutions based on reference architectures that are faster to deploy and easier to manage, that automatically react to change, and that deliver highly reliable applications for end-users, as shown in ESG’s analysis, can have a significant impact on the financial success of potential customer organizations.

Market Overview

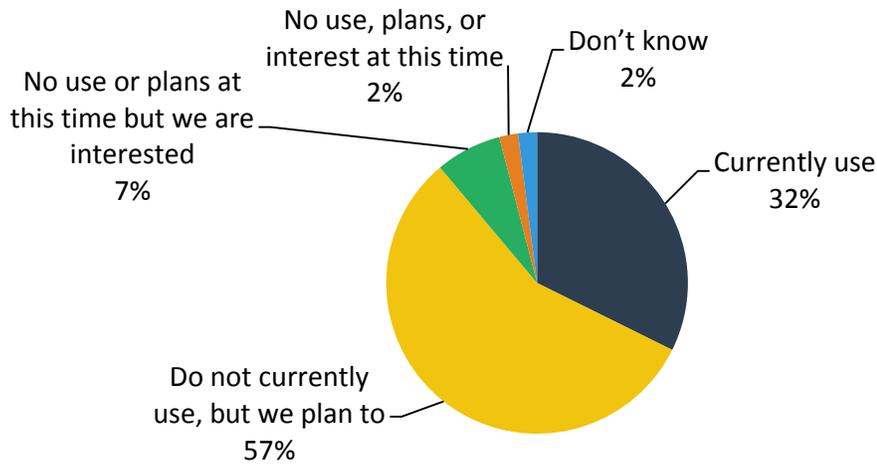
HPE has been delivering integrated and reference architecture designs for data center infrastructure since 2009. As an experienced vendor in this space, HPE has been helping IT organizations stand up data center infrastructure faster and increase the efficiency with which they deliver and maintain services to users. HPE has provided infrastructure solutions that take advantage of their world-class engineering and rigorously tested and pre-qualified configurations.

HPE’s focus in this area is well supported by market demands. In recent ESG research, ESG questioned organizations of more than 100 employees about whether they had deployed any converged infrastructure solutions. The results were striking, with nearly 9 out of 10 organizations reporting that converged solutions were in use, or planned for deployment (see Figure 1).¹

¹ Source: ESG Research Report, *The Cloud Computing Spectrum, from Private to Hybrid*, to be published.

Figure 1. Usage of Converged Computing Platforms

Please indicate your organization’s usage of or plans for converged technology solutions.
(Percent of respondents, N=308)



Source: Enterprise Strategy Group, 2016.

This same respondent pool indicated many reasons driving the adoption of converged and hyper-converged platforms, including: improved service and support (41%), improved scalability (39%), increased VM provisioning agility (37%), increased cost predictability (37%), simplified management (36%), and speed of deployment (36%), just to limit data to the top five most frequently cited responses.

The market’s perceptions and demands with respect to converged infrastructure align with the areas in which HPE’s Converged Architecture 700 platform excel:

- **Improved service and support** are delivered in two key ways. Users reliant on the applications supported by HPE’s platform are supported by a pre-qualified, -tested, and -balanced mix of servers, storage, and networking, including HPE’s ninth generation BL blade servers powered by Intel’s Xeon E5 family of processors, 3PAR StorServ storage, and customer options among HPE, Cisco, and Arista network infrastructure. Moreover, IT organizations leveraging HPE Converged Architecture 700 are provided with the proverbial “one throat to choke,” when they work with a certified HPE partner to architect, deploy, and support the platform.
- **Improved scalability** is delivered via a flexible architecture which can be right-sized for any virtual environment leveraging a building block approach in which additional servers, storage capacity, and even chassis, can be added to the configuration to support VM environments that scale from the hundreds or even thousands in a repeatable fashion.
- **Faster VM provisioning speed** is delivered via HPE’s unified management console, HPE OneView, which provides a templated approach to VM provisioning. Users can define a VM template once and reuse that template to roll out identical VM instances in a highly repeatable and automated fashion.
- **Increased cost predictability** is once again bolstered by the engineering and system design expertise of HPE. In traditional infrastructure purchasing scenarios, a customer—or even the customer’s system integrator or channel partner—is unlikely to have the expertise and resources to tune the performance and capacity capabilities of the infrastructure to achieve parity with the level of pre-qualification and testing achieved by HPE’s engineering and product development teams. The result is often a mismatch of the performance and capacity characteristics of the infrastructure and the customer’s requirements—generally resulting in the over purchasing of hardware. HPE Converged Architecture 700 customers have the assurance that the system designed by HPE and delivered by a certified partner will match their needs and eliminate cost overruns commonly associated with traditional capacity planning.

- **Simplified management** delivered by HPE OneView allows administrators to have visibility and control over each infrastructure layer in the platform and is deeply integrated with administrative tools from both VMware and Microsoft. By contrast, in typical mixed infrastructure stacks, an administrator may need to leverage multiple tools to administer each layer—from servers, to storage, to networking, to the hypervisor—hampering efficiency and increasing the opportunity for administrator error.
- **Faster speed of deployment** delivered by a vast partner ecosystem trained to deploy the platform utilizing the reference architecture provided by HPE. In traditional infrastructure purchasing scenarios, which may take several months, many vendors are typically engaged, negotiated, and transacted with individually. Additionally, the customer or their systems integrator must spend many man-hours planning and architecting the solution. Finally, the actual process of on-site integration requires expertise and one-time customization frequently fraught with difficulties.

The remainder of this paper discusses the process which ESG went through to build an economic model to quantify each of these benefits, and many others, with the intention of helping prospective customers understand the true economic impact of investing in HPE Converged Architecture 700 compared to more traditional infrastructure approaches.

HPE Converged Architecture 700: Economic Value Analysis

Methodology

For a discussion of the research and modeling methodology ESG adhered to in the process of writing this report please see Appendix A.

Economic Value Model Overview

As articulated in Appendix A, ESG's economic value analysis compares two scenarios: The first is an organization that elects to use HPE Converged Architecture 700 to support its virtual compute infrastructure requirements. The second scenario is a PMO that reflects a more conventional "component-based" approach that most customers currently take to meet their virtual infrastructure requirements. The basic profiles for each scenario follow:

- **HPE Converged Architecture 700 scenario:** In this scenario, the customer is using HPE Converged Architecture 700—a reference architecture that combines HPE servers, storage, networking, with VMware virtualization software, and tailored systems-management capabilities (delivered by HPE OneView), but which is still customizable when installed at the customer data center—via the ability to leverage existing customer equipment that can be included in the build. ESG's model takes into account the purchase price of the system components, VMware software and maintenance costs, and related IT labor costs for planning, ordering, implementing, administering the system, and training.
- **PMO scenario:** In this scenario, the customer is using a comparable set of hardware and software components that are selected, installed, and configured manually by a systems integrator on the customer's premises. ESG's analysis assumes that the customer is using blade servers and SAN storage, and that the configuration will be clustered using 10 GB Ethernet switches and networking interfaces for server interconnects within the cluster, for SAN storage, and for external access to the servers. ESG also assumed redundant network switches and interconnects for high availability. Comparable operating system, hypervisor, and systems-management software is also assumed in this scenario. ESG's model takes into account all hardware, software, and data center infrastructure costs associated with this solution, plus related IT labor costs for planning, ordering, implementing, ongoing environment administration, and training.

For both scenarios, ESG modeled the costs and IT savings benefits associated with the following tasks:

- Planning and architecture/design tasks required to scope the solution and prepare for deployment
- Deployment tasks including initial installation and setup, plus periodic upgrades and ongoing maintenance activities
- IT administration tasks such as provisioning and configuring new virtual servers and applications
- IT administration tasks related to storage and network installation, configuration, provisioning, and management
- Change management tasks performed as new software is added to virtual servers, existing software applications are upgraded, virtual machines are migrated across physical server resources, and existing infrastructure components are upgraded or replaced
- Ongoing systems management activities performed by the system and personnel for monitoring system activity, taking actions, and reporting on system status
- Note that ESG's model considers both the current infrastructure needs and the expected growth of the environment over three years to size the configuration of the solutions considered at the outset of the three-year time horizon

Simply put: ESG's analysis estimates the likely cost and potential benefits of implementing and managing—according to the tasks outlined—both HPE Converged Architecture 700 and component virtual infrastructure solutions.

Cost Categories

This ESG analysis considers six cost categories: hardware, software, infrastructure, maintenance and support, professional services, and staff costs. The sum of these categories equals the total cost of ownership (TCO) of each solution—see Appendix B for full details.

Benefit Categories

This ESG analysis considers three primary benefit categories: IT efficiency savings, user productivity improvements, and application time to value improvements delivered by the infrastructure solution selected. The sum of these categories equals the total benefit of each of the given infrastructure approach—see Appendix B for full details.

Default Scenario

ESG developed a baseline profile of a hypothetical enterprise to illustrate the relative costs and benefits of HPE Converged Architecture 700 compared with the PMO discussed in this report. For the purposes of this analysis, ESG tuned its assumptions to be representative of a growing, enterprise-sized virtual environment consisting of 500 virtual machines at the outset of the time horizon, growing by 100 VMs annually over the three-year time horizon.

To model the impact of different virtual infrastructure solutions on application environments, and ultimately, end-users, ESG also uses inputs related to the application profile of the hypothetical enterprise to calculate the number of end-users potentially affected by administrative and application availability events. ESG's model allows for three tiers of application workloads: heavy workloads—accounting for 10% of total application mix, moderate workloads—accounting for 40% of the application mix, and light workloads—accounting for the final 50% of the application mix.

At a high level, heavy workloads are intended to represent resource-intensive applications with two CPU cores and more than 150 IOPS being allocated per VM. Moreover, it is assumed on average that five VMs will be allocated per application to support this tier of workload. On average, this tier of application is assumed to be supporting 250 concurrent users per application. The moderate tier of workloads are characterized by VMs with one CPU core and 50-100 IOPS allocated. Additionally, it is assumed that this tier of application has an average of two VMs per application dedicated. On average, this tier of application is assumed to be supporting 100 concurrent users per application. Finally, light application workloads are characterized by VMs with .25 physical CPU cores and <50 IOPS allocated to them with applications and VMs existing in a 1-to-1 relationship. On average, this tier of application is assumed to be supporting 25 concurrent users per application. In total, 576 applications are assumed to be supported by the 800 VMs present in the environment at the end of the time horizon. The monthly application value for this portfolio of applications is assumed to be equal to \$850,000. For the purposes of this analysis, this application value figure includes revenue implications associated with being able to deploy applications on the infrastructure more quickly with HPE Converged Architecture 700 compared with the PMO. This assumption is important as these implications vary widely organization-to-organization and for many organizations this assumption may be conservative.

These and other key assumptions used in ESG's default scenario are summarized in Table 1.

Table 1. Key Default Scenario Assumptions for Typical Enterprise Use Case

Parameter	Default Assumption
Initial number of VMs at deployment	500
Annual growth of VMs	100
Assumed percent of applications that meet heavy/moderate/light workload profile	10% / 40% / 50%
Typical number of concurrent users for applications meeting heavy/moderate/light workload profiles	250 / 100 / 25
Average number of VMs allocated to application that meet heavy/moderate/light workload profile	5 / 2 / 1
Number of VMs per CPU core for VMs supporting applications meeting heavy/moderate/light workload profiles	.5 / 1 / 4
Average amount of storage per VM	100 GB
Average monthly value of applications supported	\$850,000
Average annual burdened cost – typical IT administrator	US\$80,000
Average annual burdened cost – typical employee (application user)	US\$65,000
Time horizon of analysis	3 years
Cost of capital	15%

Source: Enterprise Strategy Group, 2016.

Summary of Results

With the model parameters tuned to the default assumptions in Table 1, ESG’s analysis concludes that the net benefits of implementing HPE Converged Architecture 700 to support a broad portfolio of enterprise applications greatly outweigh the associated costs. Table 2 shows the modeled return on investment (ROI), project payback period, net present value (NPV), annual total cost of ownership (TCO), and annual benefit over the time horizon for an HPE Converged Architecture 700 deployment compared with a similarly-sized component-based alternative approach. The following section details the most compelling findings from this analysis as they relate to both the costs and benefits associated with HPE Converged Architecture 700 and how they differ from traditional virtual computing infrastructure approaches.

Table 2. Economic Value Summary, HPE Converged Architecture 700 versus the PMO

Scenario	Project ROI	Payback Period (years)	Net Present Value (NPV)	Annual TCO	Annual Benefit
HPE Converged Architecture 700	156%	.72	\$1,502,421	\$806,892	\$2,068,732
PMO	-27%	3.73	(\$628,092)	\$868,796	\$805,248

Source: Enterprise Strategy Group, 2016.

Annual TCO

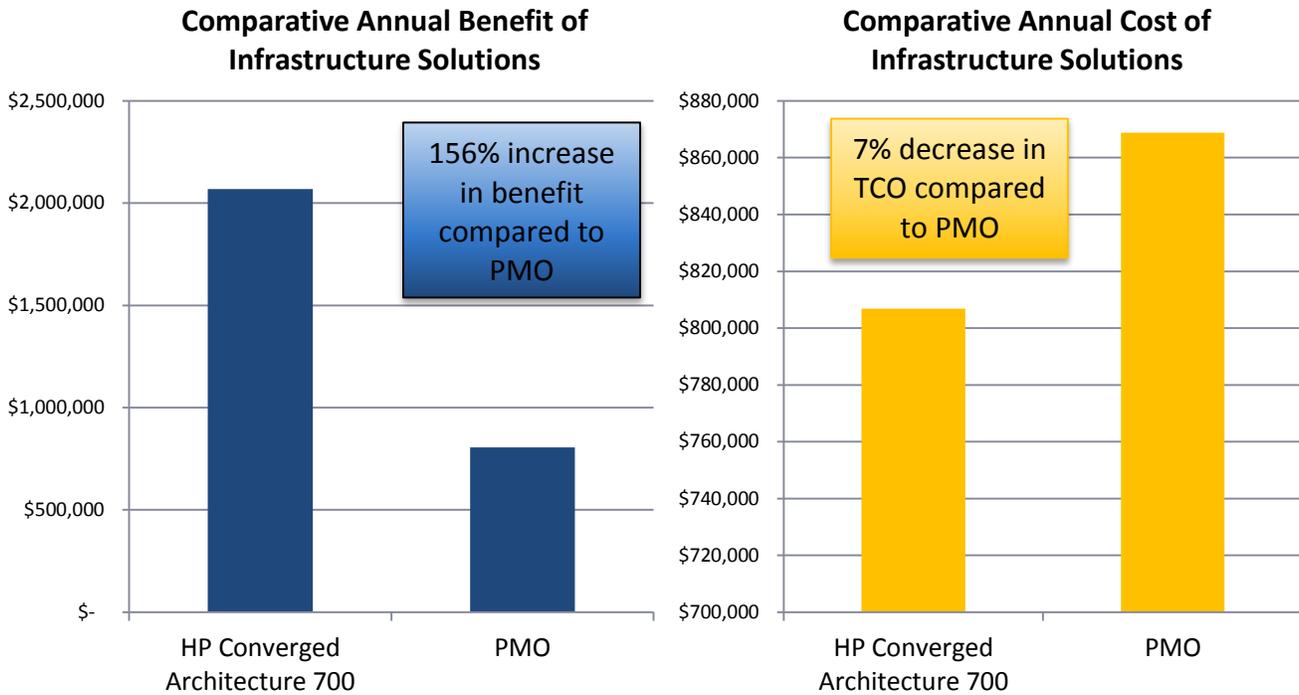
Annual TCO is the sum of all the cost categories included in the analysis (as outlined in Table 5), averaged over three years. As displayed in Table 2, the annual TCO for HPE Converged Architecture 700 is estimated as \$806,892, a moderate 7% savings compared to the PMO. However, TCO should be only one part of the customer consideration when weighing available virtual computing approaches. As shown in Table 2—and discussed in this report section—the lower costs associated with HPE Converged Architecture 700 are augmented by significant

benefits in the area(s) of increased IT efficiency, improved user productivity, and improved application time to value.

Annual Benefit

Annual benefit is the sum of all the estimated benefit categories included in this analysis (as outlined in Table 6), averaged over three years. As displayed in Table 2, the annual benefit associated with HPE Converged Architecture 700 is estimated as \$2,068,732, compared with \$805,248 for the PMO. The modeled annual costs and benefits for both scenarios is depicted graphically in Figure 2.

Figure 2. Annual TCO and Benefit, HPE Converged Architecture 700 versus the PMO



Source: Enterprise Strategy Group, 2016.

ROI

ROI is a profitability ratio for investments. It is calculated by dividing the net benefits of an investment (i.e., the total benefits minus the associated costs) by the total cost of the investment. A positive ROI indicates that total benefits exceed the costs of the investment. As displayed in Table 2, the modeled ROI for HPE Converged Architecture 700 using the inputs defined in Table 1 is 156% (a significantly higher ROI than is estimated for the PMO).

Payback Period

Payback period is an estimate of when a customer will start to see a positive return from the virtual infrastructure solution they select; it measures benefits achieved over time and costs incurred over time and indicates the investment’s break-even point. As displayed in Table 2, the expected payback period for an HPE Converged Architecture 700 deployment in an environment described by the inputs in Table 1 is .72 years or just under 9 months (significantly shorter than the payback period estimated for the PMO).

Net Present Value (NPV)

NPV is a measure which calculates the difference between the present value of cash returns and the present value of cash outflows associated with a project. It assumes a discount rate to calculate the present value of future returns. This metric is commonly used in accounting organizations to evaluate projects. Initiatives with positive NPVs are generally considered to be worthwhile investments. As displayed in Table 2, the modeled NPV for HPE Converged Architecture 700 using the inputs defined in Table 1 is in excess of \$1.5M (significantly higher than the NPV calculated for the PMO).

Benefits Analysis

Potential customers evaluating modes of virtual infrastructure deployment must be cognizant of the benefits—in this analysis, broken down into IT efficiency savings, user productivity improvements, and application time to value improvements—they will achieve from that technology solution. The three-year itemized benefits for HPE Converged Architecture 700 compared with the PMO alternative ESG developed are displayed in Table 3. As shown, total benefits for HPE Converged Architecture 700 (estimated as \$6,412,135) are modeled as more than three times that of the PMO (estimated as \$1,941,607).

Table 3. Three-year Benefits, HPE Converged Architecture 700 versus the PMO

Category	Converged Architecture 700	PMO
IT efficiency savings	\$2,071,371	\$1,021,310
<i>Initial System Design, Integration, Deployment, and Configuration</i>	<i>\$208,260</i>	<i>\$77,865</i>
<i>Systems Maintenance, Support, and Management</i>	<i>\$401,444</i>	<i>\$155,815</i>
<i>Resource Management Including Storage and Network Management</i>	<i>\$895,313</i>	<i>\$534,375</i>
<i>VM Administration</i>	<i>\$566,354</i>	<i>\$253,255</i>
User productivity improvements	\$2,859,825	\$1,394,433
<i>Application Deployment</i>	<i>\$916,383</i>	<i>\$500,235</i>
<i>Application Support/Management</i>	<i>\$1,100,731</i>	<i>\$471,385</i>
<i>Reduction Planned and Unplanned Downtime</i>	<i>\$842,711</i>	<i>\$422,813</i>
Application time to value improvements	\$1,275,000	\$0
Total three-year benefits	\$6,206,196	\$2,415,743

Source: Enterprise Strategy Group, 2016.

Major Benefit Differences for HPE Converged Architecture 700 versus the Built-by-customer PMO

Benefits were calculated based on observations and estimates related to the value of HPE Converged Architecture 700 and the PMO obtained through past ESG quantitative and qualitative research, relevant product demos, literature reviews, and in-depth interviews with technical stakeholders at HPE.

Key IT efficiency benefit assumptions for HPE Converged Architecture 700 and the PMO are:

- Initial System Design, Integration, Deployment, and Configuration:** In a component-based IT infrastructure, the burden of solution design and integration falls on the IT organization and their selected systems integrators. With HPE Converged Architecture 700, much of the planning, system balancing, and rigorous pre-testing for various workloads resides with HPE’s engineering resources. Moreover, HPE’s certified partners conduct an on-site product orientation and training with each customer purchasing an

HPE Converged Architecture 700. These advantages dramatically decrease the deployment time and effort for the customer.

To capture these differences, by default, ESG's model assumes eight hours of IT staff time will be dedicated to architecture and planning operations in an HPE Converged Architecture scenario compared to in excess of 50 hours in the PMO scenario. Actual IT staff time dedicated to system deployment in the HPE Converged Architecture deployment is modeled to require only six hours of IT staff time compared to the PMO's 50. Additionally, ESG's model assumes material differences in the time and effort required to set up the management tools for the infrastructure. With HPE Converged Architecture 700, HPE OneView is the only tool needed to manage the entire stack. As such, ESG's model assumes about a one hour set up and configuration process. Considering a piece-part infrastructure stack consisting of potentially many vendors' products, each with their own native management tools, ESG's model assumes a much more labor-intensive 20 hours of management tool setup and configuration tasks.

Each of these project tasks, in addition to marginal assumed improvements in the assumed time to apply regular platform firmware updates (60 minutes per assumed biannual update with HPE vs. 8 hours per major update with the PMO) and storage configurations (10 minutes to configure storage for each virtual server added to the environment over time vs. 1 hour) are the major contributors to the incremental \$131,239 in IT efficiency estimated for the HPE scenario compared to the PMO.

- **Systems Maintenance, Support, and Management:** Beyond the setup and configuration efficiencies associated with a reference architecture approach, which are weighted towards the beginning of the time horizon, a number of IT efficiencies are created over time. Leveraging a pre-tested and professionally-installed reference architecture compared to completely custom solutions is estimated in ESG's model to ease the ability to add physical servers to the platform over time as requirements change, reduce the number of server configuration errors and reworking, and reduce the number of support calls to the partner providing support (vs. potentially multiple vendors) over time.

To account for the fact that in the HPE scenario, the customer is reliant upon a pre-validated reference architecture installed by a knowledgeable partner, physical server additions are assumed to be much simpler operations in the HPE scenario. While each physical server is modeled to require 60 minutes of IT staff time to deploy in the HPE scenario, 23 man-hours are assumed to be allocated to each physical server addition in the PMO scenario.

Similarly, server configuration errors—and the difficulty to resolve them—are modeled to be dramatically reduced in the HPE scenario. For each physical server added to the environment over time, ESG's model assumes that 10% will require some degree of rework. However, to capture the facts that this frequency may in fact be lower and issues will certainly be easier to resolve in the HPE scenario, the average resolution time in the HPE scenario is assumed to be 30 minutes vs. nearly three work days—essentially equivalent to a redeployment—in the PMO scenario.

Finally, reflecting the fact that, in the HPE scenario, customers have one throat to choke in the event that something goes wrong—an HPE partner who only has to interface with one vendor—ESG's model assumes that customers will spend less time working with their HPE VAR to resolve those issues compared to the PMO. In the PMO scenario, a customer, or their VAR, may need to reach out to many different vendors, each of which may shift the responsibility of problem resolution on to other vendors in the stack. To account for this difference, ESG's model assumes a 20% increase in the number of support calls made by IT in the HPE scenario vs. the PMO (5 annually vs. 6.25) and assumes a 1.5 hour increase in IT efficiency related to issue resolution (6 hours per call vs. 7.5) in the HPE scenario.

These events, and their respective IT labor cost per-event, are the major drivers behind the incremental \$245,629 of IT efficiency observed for the HPE scenario compared to the PMO.

- **Resource (Storage and Network) Management:** Although the configuration and provisioning of storage for VMs over time is accounted for in the "initial system design, integration, deployment, and configuration" IT efficiency line item, there are many physical storage and networking administration tasks modeled in the

environment over time which are eased and improved thanks to HPE's ability to manage infrastructure holistically, via HPE OneView, and the higher degree of component integration achievable in a single-vendor reference architecture compared to a piece-part infrastructure. These benefits are captured in the "resource management" efficiency area.

Included in the resource management efficiency line item are physical storage and networking set up tasks, which are estimated on a per-new VM basis over time. In the HPE scenario, an estimated 20 minutes of IT staff time is dedicated to networking and SAN configurations compared to 1.25 hours per VM in the PMO scenario for each VM added to the environment. Additionally, to account for steady state tuning and resource balancing over time, for each VM in the environment, a biannual, ten-minute labor cost is allocated to LAN and SAN moves, adds, and changes. This compares to an estimated biannual, forty-minute labor cost in the PMO. In the aggregate, over three years, these assumed differences result in an incremental \$360,938 of IT efficiency in the HPE scenario compared to the PMO.

- **VM Administration:** HPE OneView not only eases the lives of IT administrators with respect to physical resource provisioning and administration, but its deep integration with VMware vSphere is also modeled to create material efficiencies in the realm of VM administration tasks. In particular, the ability to create virtual server templates and provision VMs from those templates in an automated and error-free manner is modeled to be markedly improved compared to more manual and error prone VM provisioning which is typical of traditional piece-part infrastructures.

To account for the advantages of HPE OneView, ESG's model first assumes a marginal improvement in the effort required to provision each virtual server in the environment—15 minutes in the HPE scenarios compared to 25 minutes in the PMO. Additionally, to account for fewer errors occurring during VM setup, ESG's model assumes a nearly 50 times multiplier between the IT staff time dedicated to resolving VM provisioning errors in the PMO compared to the HPE scenario. These two assumptions lead to the \$313,099 improvement in VM administration efficiency estimated in the model.

The HPE Converged Architecture clearly creates a number of efficiencies which impact the IT organization. However, many of these efficiencies also trickle down to the user communities of the various applications supported by the infrastructure. ESG's model also quantifies the value of these user improvements in terms of expected gains in productivity based on the average fully burdened salary of application users. However, it is important to note that ESG's model aims to be conservative, only counting 80% of time saved for users as productive. Key user productivity benefit assumptions for HPE Converged Architecture 700 and the PMO are:

- **Application Deployment:** As noted in detail, in an HPE use case both virtual server provisioning and reconfigurations are estimated to be completed much more efficiently. While significant, VMs do not run in a vacuum. They support the applications end-users rely upon to do their jobs. Any delays in provisioning time result in periods of application unavailability and lost productivity. ESG's model accounts for this fact by allocating relative application setup times of 20 minutes for HPE and 45 minutes for the PMO, scaled by the number of total applications in the environment (576 in ESG's default scenario) and the average number of application users (45 per application in ESG's default scenario). The total difference in productive time is then multiplied by the average fully burdened productivity of end-users and adjusted down by 20%. The result over three years of operations is that the HPE scenario creates an estimated \$267,786 of incremental user productivity over the time horizon related to application provisioning and deployment.
- **Application Support/Management:** Beyond application deployment operations, ESG's model includes estimates related to ongoing application administration tasks that may impact availability, and thus, the productivity of users. For every application in the environment, ESG's model estimates biannual patch events, as well as biannual move, add, and change events. ESG's model assumes a near 50% improvement in these tasks. Thus there is an observed 50% improvement in the availability of applications across the entire application user environment during those events, resulting in an incremental \$420,306 in user productivity in the HPE scenario.

- **Reduction Planned and Unplanned Downtime:** A pre-qualified and tested platform such as HPE Converged Architecture 700 is much less likely to have interoperability problems compared to the PMO and is therefore estimated in ESG’s model to be more reliable over time. This benefit, coupled with the error prevention features such as automated resource provisioning and proactive monitoring capabilities, has a profound impact on application end-users in terms of the planned and unplanned downtime estimated in the environment.

ESG’s model accounts for these differences by making conservative assumptions about the frequency of planned downtime (1 event annually in the PMO, with a 10% reduction in events in the HPE scenario), the duration of planned downtime (10.5 minutes of lost productivity among all application users in the PMO versus 7 minutes in the HPE scenario), the frequency of unplanned downtime (25% chance of an event in the PMO, with a 10% reduction in the likelihood of an event in the HPE scenario), and the duration of unplanned downtime events (35 minutes of lost productivity among all application users in the PMO versus 30 minutes in the HPE scenario). The result is a total delta in expected user productivity of \$281,532, favoring the HPE scenario.

The final benefit area HPE Converged Architecture 700 is modeled to deliver in the scenario examined in this paper is related to application time to value. This value is attributed to the fact that a converged reference architecture provides a blueprint for a VAR to deploy. With custom, one-off infrastructure deployments integrating many different vendors, the elapsed time to stand up the infrastructure and make it “production ready” is assumed to be much longer. ESG’s model leverages the assumed application value of \$700,000 and an assumed deployment time of 3 months for the PMO versus 1.5 months in the HPE scenario to estimate a total of \$1,050,000 of incremental application value created for the hypothetical organization described in this report.

TCO Analysis

For the hypothetical customer scenario described in Table 1, the estimated three-year TCO for HPE Converged Architecture 700—compared with the TCO estimated for the PMO—is displayed in Table 4. As shown, from a TCO perspective, HPE Converged Architecture 700 is expected to be slightly less expensive than the PMO over a three-year time horizon.

Key TCO assumptions:

- It is important to note that in ESG’s model, since the virtual environment to be supported is identical whether the organization is leveraging HPE Converged Architecture 700 or the PMO, the hardware and software costs incurred by the customer are assumed to be identical. ESG believes this to be a very conservative assumption, as it is possible that HPE customers may be able to effectively “get more out of” their infrastructure investments due to the fact that HPE Converged Architecture 700 has the advantage of being a balanced reference architecture with the weight of HPE’s engineering expertise behind it. However, this model is focused on the benefits associated with converged versus traditional infrastructure deployment models and for the sake of clarity ESG did not incorporate any assumed hardware or software savings in its model by default.
- It is also important to note that while HPE Converged Architecture 700 is a converged platform, it is unlike other rigidly structured converged systems which do not allow for easy customer customization or allow the customer to reuse past infrastructure investments. **HPE Converged Architecture 700 is flexible, in that, if customers have excess conforming server, switching, or storage infrastructure in their environment today, they can redeploy that in their converged platform.** In this particular scenario, the customer is assumed to have half of the required blade servers needed to support the virtual environment on hand at the start of the time horizon. While this does not drive a differential with the PMO, as the same reduction in TCO could be argued in that scenario, it does increase the ROI for the solution as the TCO is assumed to be reduced by slightly more than \$91,000 relative to the benefits observed.

Table 4. Three-year TCO, HPE Converged Architecture 700 versus the PMO

Category	HPE Converged Architecture 700	Built-by-customer PMO
Hardware	\$1,285,998	\$1,285,998
Software	\$293,758	\$293,758
Infrastructure	\$11,500	\$11,500
Maintenance & Support	\$292,253	\$292,253
Professional Services	\$501,288	\$558,916
Staff Personnel	\$35,881	\$163,926
Total three-year costs	\$2,420,677	\$2,606,387

Source: Enterprise Strategy Group, 2016.

Major Cost Differences for Converged Architecture 700 and the PMO

As noted, hardware and software costs are identical in both the HPE scenario and the PMO. Similarly, data center infrastructure and maintenance and support costs between the two scenarios, which are correlated to hardware and software spend, are also normalized. However, the simpler to architect, deploy, and manage HPE scenario is assumed to yield material savings in the areas of professional services costs and internal staff personnel costs.

- Professional Services:** Professional services costs are incurred in both scenarios and are made up of three components: pre-sales consulting, installation services, and ongoing support.

Presales support consists of all out of pocket expenses paid to a VAR during the planning and architecting phase of the project. For the VAR, efficiencies are gained by leveraging a reference architecture validated by HPE. This is compared with a scenario where multiple vendors' products must be procured and architected individually. ESG's model assumes that some of these efficiencies will be passed on to the customer in the form of lower costs. ESG's model assumes that the total expenditure for presales support in an HPE Converged Architecture 700 is 5% of the total system hardware and software CapEx, or \$83,548. By contrast, ESG's model assumes a 10% increase, or \$91,903, in the PMO scenario.

Installation services consist of all the out of pocket expenses associated with installing and getting the platform ready for production on the customer premises. Similarly, this process is assumed to be more efficient for the VAR when leveraging the HPE reference architecture and savings are once again assumed to be passed on to the customer. For HPE installation services, ESG's model uses the best practice parameter of 10% of the system cost, or \$167,096. This compares favorably to the \$183,805 estimated in the PMO scenario.

Finally, an ongoing support contract is assumed to be entered into by the customer in both the PMO and HPE scenarios. Again, the validated and more stable architecture provided by HPE is assumed to make supporting the platform over time markedly easier for the VAR and thus allow the VAR to offer the customer a less costly contract rate. ESG's model assumes the customer will spend \$250,644 over the three years for ongoing support services compared to \$275,708 in the PMO scenario.

- Staff Personnel:** While many staff time improvements for tasks are captured in ESG's model by the IT efficiency savings area of the model, there are additional, less operational, savings to capture. Time spent focused on project management, system/solution procurement, and meetings to discuss solution planning and operations, which do not directly impact day-to-day operation of the platform over time are captured in the staff personnel cost category. The added simplicity of consuming a single-vendor reference architecture reduces these costs as compared with traditional IT purchasing scenarios. In total, over three years, ESG's model estimates savings in this area to be \$128,082, or slightly more than one fully burdened IT FTE.

The Bigger Truth

As evidenced in ESG research, IT decision makers are broadly adopting converged infrastructure solutions. Moreover, IT organizations are doing so with a priority placed on the environment agility and service excellence, both for themselves (i.e., a stable environment with leading support from VARs and vendors) and for their users (i.e., the ability to make their end-users happy). ESG research also shows that CapEx and OpEx considerations take a backseat to these more strategic desired outcomes with respondents citing these areas as adoption drivers at frequencies of 26% apiece.²

However, ESG's research and financial modeling conducted for the HPE Converged Architecture 700 show that dramatic business value, in the form of IT efficiency, user productivity, and application time to value, can be delivered by this converged platform, while allowing for the achievement of TCO parity, or even modest gains, with traditional infrastructure deployment models.

For IT organizations looking to increase their efficiency, while boosting the level of service delivered to application users in a cost-neutral manner, HPE Converged Architecture 700 warrants close consideration.

² Source: ESG Research Report, *The Cloud Computing Spectrum, from Private to Hybrid*, to be published.

Appendix A

For this project, ESG adhered to the following research and modeling methodology:

- ESG conducted initial market research across HPE and other relevant IT vendors to assess current market trends, vendor value claims, and the purchase considerations that are most important and relevant to existing and prospective converged computing infrastructure customers.
- Based on the results of this initial research, ESG subsequently:
 - Identified a “present mode of operation” or PMO—effectively, the alternative approach that customers are likely to take to meet their virtual infrastructure requirements—against which the costs and benefits of utilizing HPE Converged Architecture 700 was to be compared. A conventional component-based infrastructure based on a combination of individually selected, tested, and integrated compute, storage, and network products was used.
- ESG then developed a comprehensive financial model designed to qualify and quantify the potential costs and benefits of utilizing HPE Converged Architecture 700 compared with the PMO.
- ESG then conducted a series of in-depth interviews with systems engineering, service and support, and technical marketing representatives from HPE. The data collected in these interviews was used to refine assumptions built into the model related to current customer environments and the direct and indirect costs and benefits attributable to both HPE Converged Architecture 700 and alternative virtual computing delivery approaches. Product demonstrations, configurators, and case studies of HPE Converged Architecture 700 customers were also used to identify specific IT tasks and the labor burden (in both time and cost) associated with those tasks. These findings were then compared against the results of ESG’s qualitative and quantitative market research with organizations currently using integrated systems (inclusive of HPE and other vendors). This research helped to inform ESG’s understanding and analysis of integrated computing adoption drivers, usage trends, and the technical, operational, and financial benefits that have been realized by customers.
- Once the economic model was finalized and all validation was complete, ESG modeled a default scenario that was designed to demonstrate the relative costs and benefits of HPE Converged Architecture 700 in a representative enterprise environment. Those results were then compared with model outcomes for a similar-scale traditional virtualization solution based on separate compute, storage, and network components. The results for this default scenario are described in the body of this paper.

Please note that the data and conclusions presented in this report regarding the costs and benefits associated with implementing HPE Converged Architecture 700 compared with alternative virtualization solutions reflect the output of ESG’s economic value analysis based on the specific use case and default scenario assumptions modeled for this report. ESG acknowledges that changes to these assumptions will lead to a different set of results and, as such, advises IT professionals to use this report as one validation point in a comprehensive financial analysis process prior to making a purchase decision. HPE provided current standard pricing and configuration information for HPE Converged Architecture 700 to ESG. Other IT equipment and labor cost assumptions were obtained from publicly available sources such as IT vendor and channel partner websites and published price lists.

Appendix B

Table 5. Cost Categories in the Scope of the Analysis

Category	Description
Hardware	<ul style="list-style-type: none"> • For both HPE Converged Architecture 700 and the PMO, this includes capital costs associated with enclosures, management servers, blade servers, network switches, and storage. <ul style="list-style-type: none"> ○ Note that ESG utilized hardware CapEx costs largely normalized such that the infrastructure platforms would be truly comparable from a performance and capacity perspective.
Software	<ul style="list-style-type: none"> • For HPE Converged Architecture 700: <ul style="list-style-type: none"> ○ Licenses for VMware vSphere. ○ Licensing costs associated with HPE OneView system management software. • For the built-by-customer PMO : <ul style="list-style-type: none"> ○ Licenses for VMware vSphere. ○ Licenses for comparable systems management and administrative tools. • Note that application stack components, such as SQL Server, web servers, or other components, are not included in the scope of the analysis. ESG made the general assumption that software CapEx for both vSphere and alternative system management software would be normalized. However the level of integration and capabilities of OneView are assumed to be superior to the PMO, reflecting the likelihood of multiple alternative tools being used to manage the PMO infrastructure versus a holistic management platform.
Infrastructure	<ul style="list-style-type: none"> • For both scenarios, this includes costs related to data center infrastructure, such as racks, cables, power, cooling, and overhead. Since both infrastructure builds are assumed to be identical in terms of their hardware footprint, no differences between HPE Converged Architecture 700 and the PMO’s infrastructure cost are estimated.
Professional services	<ul style="list-style-type: none"> • For both scenarios, this includes additional services from IT suppliers and/or third-party firms to provide planning, design, implementation, best practice consulting, and support services. These include both one-time activities at the time of solution planning and installation as well as periodic additional activities (such as training) on an ongoing basis.
Staff	<ul style="list-style-type: none"> • For both scenarios, this includes IT staffing costs associated with tasks such as project management, system/solution procurement, and meetings to discuss solution planning and operations. Additional outcomes of reducing the time associated with solution administration are captured as benefits. • All position-specific salary information is based on publicly available median average salary data. For all employees, the model uses a fully burdened rate (i.e., including the cost of benefits, payroll taxes, etc.) of 40%.
Maintenance and support	<ul style="list-style-type: none"> • For both the HPE Converged Architecture 700 scenario and the PMO, support costs include both support for VMware software components and support for the solution components. Non-VMware support costs for HPE Converged Architecture 700 are taken from current HPE pricing while those costs for the alternative integrated system are estimated to be comparable with those from HPE.

Source: Enterprise Strategy Group, 2016.

Table 6. Benefit Categories in the Scope of the Analysis

Category	Description
IT efficiency savings	<ul style="list-style-type: none"> • For both scenarios, this includes operational cost savings for IT staff related to the following activities: <ul style="list-style-type: none"> ○ Initial system design, integration, deployment, and configuration ○ Systems maintenance, support, and management ○ Resource management including storage and network management ○ VM administration • These savings are calculated based on the estimated number of IT staff hours required to perform those tasks above compared with a baseline.
User productivity improvements	<ul style="list-style-type: none"> • For both scenarios, this includes operational cost savings for application users related to the following activities: <ul style="list-style-type: none"> ○ Improved end-user productivity that results from faster virtual server and application installation and configuration ○ Improved end-user productivity that results from faster and less-disruptive virtual server and application upgrades and changes ○ Improved end-user productivity that results from reduced application downtime and performance impacts • These savings are calculated based on the estimated number of application end-user hours that are positively or negatively impacted by these events. For these calculations, ESG also assumes a productivity correction factor of 0.8 (i.e., the model assumes that 80% of the end-user hours “saved” will actually be productive hours).
Application time to value improvements	<ul style="list-style-type: none"> • For both scenarios, this includes the estimated value achieved by the organization from having applications available to users faster. ESG’s model uses a conservative value of all applications in the environment of \$700,000. This value is meant to be representative of the value of having the applications themselves available, rather than a measure of the productivity improvements of the user community.

Source: Enterprise Strategy Group, 2016.



Enterprise Strategy Group | **Getting to the bigger truth.**

20 Asylum Street | Milford, MA 01757 | Tel: 508.482.0188 Fax: 508.482.0218 | www.esg-global.com