When Comparing Cloud Alternatives, For the Best TCO Leverage VMware Cloud Foundation

APRIL 2019

THE NET NET

In this paper we examine the relative costs and other advantages of four different cloud infrastructure approaches, two based on private or on-premises clouds and two on public clouds. These public and private approaches can in turn be combined to create a hybrid cloud deployment. The objective is to enable businesses to evaluate which cloud approach makes the most sense for them, based on differences in TCO and other relevant factors.

Public clouds are here to stay, given their large and growing adoption by businesses and consumers alike. Now well over a decade since AWS first launched its infrastructure-as-a-service offerings, public clouds have become a popular deployment choice for both new and legacy business applications. Based on Taneja Group research, nearly every business is now running at least some of its use cases and applications in one or more public clouds. Clouds offer customers greater agility and near-infinite scalability, in addition to a flexible pay-as-you-go consumption model.

However, a large majority of businesses have decided they cannot rely on public clouds alone to satisfy their IT needs. Instead, they see hybrid clouds as a better architectural choice, enabling them to realize all the advantages of a public cloud along with broader use case support and a more flexible deployment model. More than two-thirds of IT professionals who participated in two recent Taneja Group research studies favor hybrid clouds as their long-term architecture.

For the on-premises or private cloud component of a hybrid cloud, the majority of users are starting with VMware technology and typically use two different approaches: a traditional, integrated 3-tier architecture commonly called Converged Infrastructure (CI); or a fully software-defined approach based on Hyperconverged Infrastructure (HCI). The 3-tier, CI approach utilizes loosely integrated compute, storage and networking resources, while the easiest and most comprehensive approach is based on VMware Cloud Foundation, a software-defined data center platform. Our analysis demonstrates that the software-defined VMware Cloud Foundation approach provides a simpler, more cost-effective approach to building on-premises or private cloud infrastructure.

Looking to the public cloud, businesses have a choice of whether to move all or just a subset of their on-premises workloads to the public cloud, and either run them there permanently or in hybrid fashion. We have analyzed the relative costs and advantages of two major ways to migrate and run workloads in the public cloud: moving on-premises workloads to a native public cloud infrastructure, such as native Amazon Web Services, Microsoft Azure or Google Cloud Platform; or moving them to a VMware Cloud Foundation-based public cloud, such as VMware Cloud on AWS or VMware Cloud Foundation offered as a service by one of the VMware Cloud Provider Program (VCPP) partners. As we'll see, moving to a native public cloud infrastructure requires often significant upfront refactoring and migration effort, which gives the path to a VMware Cloud Foundation-based public cloud a major cost advantage.

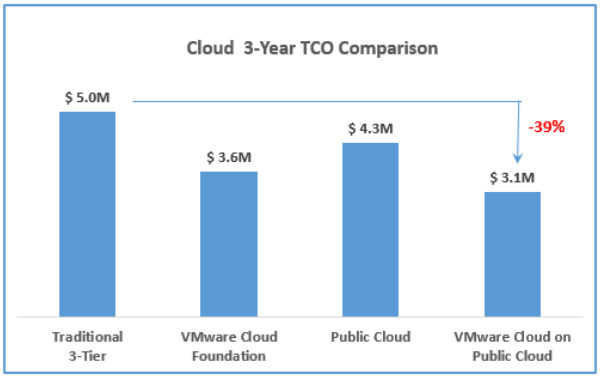

PO Box 117 - West Dennis, MA 02670 - T: 508.398.0075 - F: 508.398.0071

Based on our in-depth costing and qualitative analysis of the two private and two public cloud approaches, we found that clouds based on VMware Cloud Foundation technology offer the lowest TCO over a three-year period. VMware Cloud Foundation-based clouds minimize risk by starting with proven and widely deployed VMware technology on premises and enabling full application compatibility and workload portability between your on-premises environment and your choice of one or more VMware-compatible public clouds. VMware Cloud Foundation-enabled clouds will help you to optimize your path to a hybrid cloud deployment.

HYBRID CLOUD IS THE PREFERRED CHOICE

In recent (Nov-2018) Taneja Group research, we asked a broad cross-section of IT decision makers and influencers what IT architecture they would like to see their company adopt in the next 3 to 5 years. Faced with a choice between hybrid cloud, on-premises only, public cloud only, and colocation alternatives, respondents overwhelmingly selected Hybrid Cloud as their data center architecture of choice. Forty-five percent are already "all-in" on a hybrid cloud architectural approach, while another 20+% see colocation facilities as a stepping stone to or the equivalent of a public and ultimately hybrid cloud deployment. Only 7% plan to move fully to public cloud, and just 16% plan to remain mostly on-premises.

Why are hybrid clouds so attractive? IT professionals cite several major reasons, including the ability to cost-effectively scale infrastructure up or down as needed, transparently move workloads between on-premises and a public cloud, and dynamically place workloads in the cloud where they best fit or deliver the most value. Hybrid clouds broaden the workload deployment choice to more than one cloud and enable some compelling use cases, such as cloud-enabled backup, disaster recovery and analytics. An increasing number of organizations are also looking to lift and shift their business apps and run primary app storage in a hybrid cloud environment, as shown in Figure 1. Done right, a hybrid cloud improves IT agility while reducing overall cost.

Figure 1: Projected Hybrid/Multi Cloud Use Cases for 2019

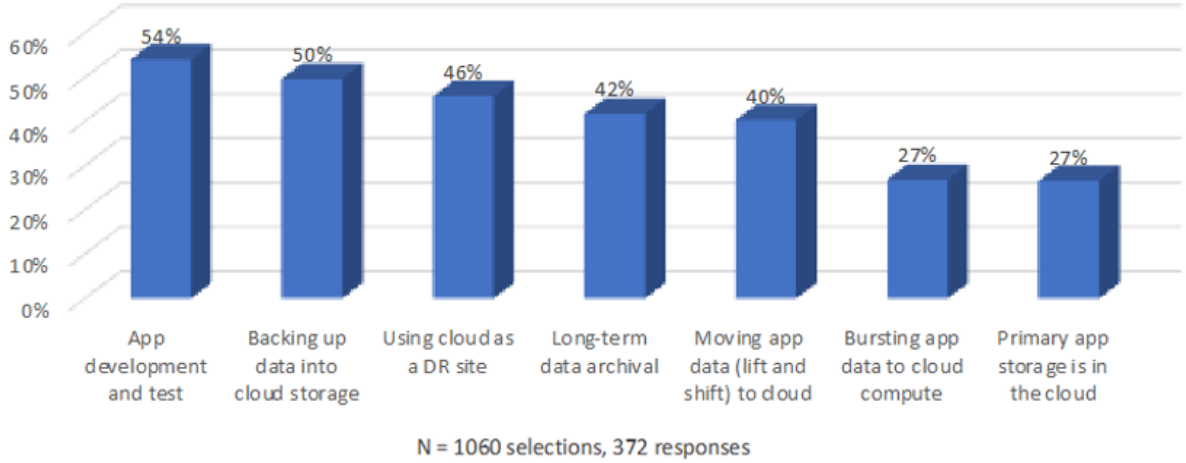

SOURCE: TANEJA GROUP

When it comes to the starting point for your hybrid cloud journey, VMware is the logical on-premises choice. A majority of enterprises (65%+) today have adopted VMware vSphere as one of their primary virtualization technologies, and tend to run many if not most of their business-critical apps on the platform. A growing percentage of VMware users are building private and hybrid clouds on VMware Cloud Foundation, the company's software-defined data center (SDDC) platform, which provides the same architecture as VMware Cloud on AWS. These customers benefit from a proven, enterprise-caliber hybrid cloud solution that provides full application compatibility and workload portability between on-premises and public cloud environments.

PO Box 117 - West Dennis, MA 02670 -T: 508.398.0075 - F: 508.398.0071

In our TCO assessment, we evaluated three major cost components of a hybrid cloud solution:

- **Private Cloud architectural approach.** In this category we looked at the procurement and operational costs of two different approaches to building a private cloud, based respectively on a fully software-defined approach and an integrated 3-tier architecture.
- **Public Clouds associated with Hybrid Cloud.** Here we analyzed the service and support costs of running on two different public clouds, VMware Cloud on AWS and Microsoft Azure.
- **The refactoring and migration cost.** Given their magnitude and one-time nature, we broke out into a separate category the costs of adapting or reengineering and moving apps from on-premises to a public cloud, looking at both packaged and custom-built applications.

We describe and quantify these hybrid cloud cost components in detail in the next few sections.

CONTRASTING THE TWO PRIMARY APPROACHES TO BUILDING A PRIVATE CLOUD

In this section of our TCO analysis we are focusing on on-premises private clouds. In this study, we are comparing two popular private cloud approaches: one built using integrated 3-tier architectures called Converged Infrastructure (CI); and another built using a fully software-defined approach commonly called a Hyperconverged Infrastructure (HCI). In the study, we will contrast each of these architectural options and evaluate the total costs to maintain the core private cloud infrastructure over three years. The following sections summarize our analysis of each of these approaches. Most businesses will likely be upgrading from some previous level of on-premises infrastructure, but in our study we felt it best to use a greenfield scenario to facilitate a level playing field for our head-to-head comparison.

VMware Cloud Foundation based Private Cloud

VMware Cloud Foundation software establishes a unified software-defined data center (SDDC) platform for both private and hybrid clouds. In the context of private cloud, customers can deploy VMware Cloud Foundation software on top of any qualified vSAN ReadyNode or integrated systems from numerous hardware vendors. VMware Cloud Foundation has been designed from the start to be the simplest path to an all software-defined private cloud. Figure 2 depicts the key components that make up an SDDC cloud based on VMware Cloud Foundation.

Figure 2: VMware Cloud Foundation Overview

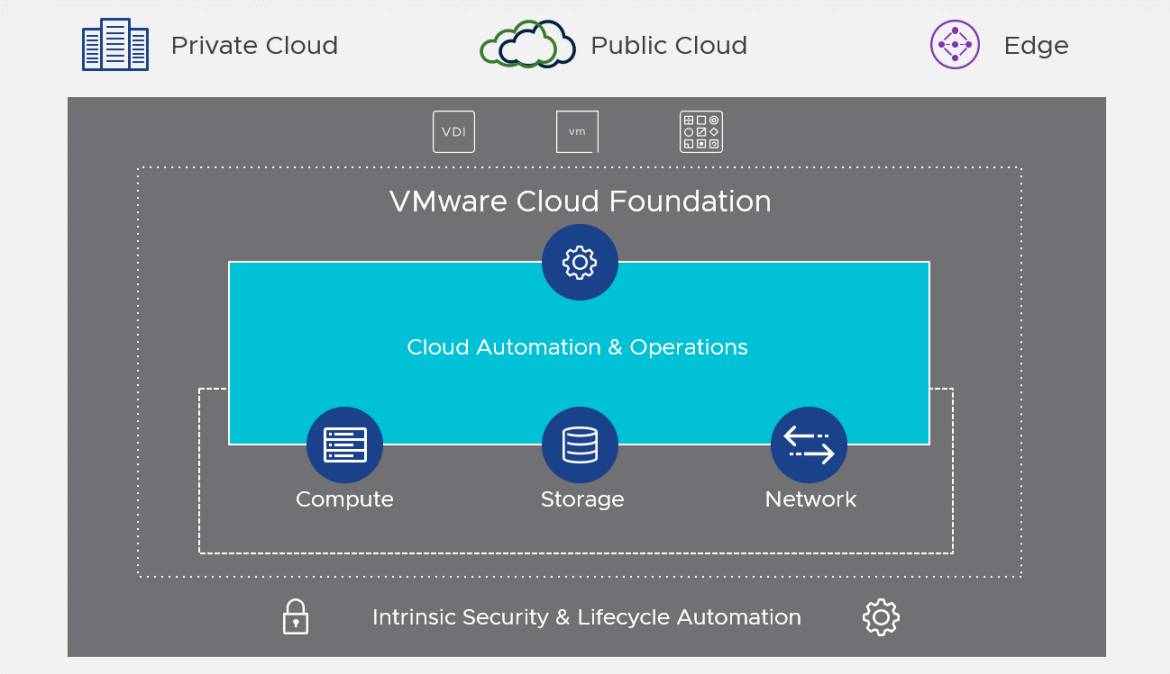

SOURCE: VMWARE

VMware Cloud Foundation integrates logical and physical compute, storage, and networking into a unified solution. It brings together VMware's compute, storage, and network virtualization into a natively integrated stack that combines hyper-converged software (vSphere plus vSAN) with network virtualization (NSX). A typical deployment contains x86-based Hyperconverged (HC) ReadyNodes with Direct Attached Storage (DAS), top of rack (ToR) switches, and Power Distribution Units (PDU). Additional racks may contain optional spine switches to interconnect racks in a highly resilient, scale-out, spine-leaf architecture. Figure 3 depicts a typical layout of a single rack.

- **Networking:** Spine, Leaf, and management switches are provided by any supported network provider of choice. VMware NSX provides network virtualization and security.
- **Management Cluster:** The management cluster uses 4 hyper-converged servers and hosts the Cloud Foundation management software and necessary components.
- **Workload Clusters:** The hyper-converged servers can be any vSAN ReadyNode provided by numerous qualified server vendors.

Figure 3: vSAN ReadyNode and Networking

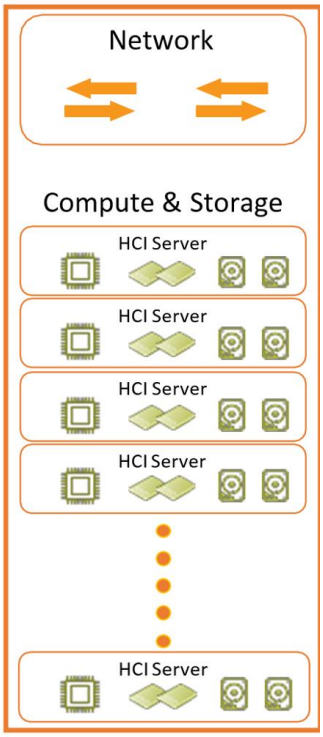

The key difference from 3-tier partial SDDC architectures is that VMware Cloud Foundation includes cloud automation tools and completely virtualizes every infrastructure layer – networking, compute, storage – using a natively integrated SDDC stack that embeds all layers as virtualized services directly within the vSphere hypervisor kernel. The VMware Cloud Foundation SDDC platform can be optimized for various IaaS workloads. Finally, because the unified and natively integrated SDDC stack can be run as a service from the public cloud, VMware Cloud Foundation enables powerful hybrid cloud capabilities.

VMWARE CLOUD FOUNDATION INTEGRATED INFRASTRUCTURE MANAGEMENT

VMware Cloud Foundation includes a tool for the cloud operator to automate Day 0 to Day 2 workflows of the entire unified SDDC software stack. VMware Cloud Foundation automates the bring-up and configuration process of the Cloud Foundation stack, including VMware infrastructure VM deployment, management cluster creation, storage configuration, cluster creation and provisioning, and more. Next, VMware Cloud Foundation enables one-click patching/upgrading of the complete software stack. Cloud admins have the flexibility to choose the timing and scope of the updates and can apply updates/upgrades on a per workload domain basis. Cloud Foundation simplifies resource allocation to individual workloads by automating cluster creation through policy-based provisioning.

In summary, VMware Cloud Foundation provides:

- Automated, simplified and rapid setup of the entire SDDC-based private cloud.
- Automated lifecycle management of the entire SDDC Platform including day 0 to day 2 processes such as bring-up, configuration, provisioning, and patching/upgrades.
- Workload Domain abstraction for isolating pools of resources into private cloud capacity with different availability, performance, and security attributes.
- Integrated management of servers and virtualized resources from a single pane.
- Operational simplicity and automation for health monitoring of both physical and virtual infrastructures.
- Simplified IaaS service deployment and operations similar to a public cloud.

As a result, the VMware Cloud Foundation solution significantly improves agility and reduces TCO, enabling a simpler and cheaper way to deliver an SDDC private cloud.

PO Box 117 - West Dennis, MA 02670 -T: 508.398.0075 - F: 508.398.0071

Traditional 3-Tier Architecture based Private Cloud

Another popular private cloud approach involves using traditional 3-tier systems with several software components added in order to achieve a partially integrated SDDC solution. For this study, we analyzed a leading traditional 3-tier integrated system (or converged infrastructure, CI). We based our evaluation on external shared all-flash storage and blade-based computing infrastructure along with the appropriate VMware virtualization technology needed to complete the solution. Figure 4 shows the basic approach to the rack design using a traditional 3-tier architecture.

- **Networking:** Spine and Leaf based on networking provided by industry leaders. NSX software-based network virtualization was used to minimize cost and align functionality between the solutions we analyzed.
- **Management Cluster:** A dedicated set of isolated servers forms a compute cluster which is sized to host all dedicated management software for operating both the hardware and virtualization software infrastructure.
- **Compute Cluster:** A dedicated set of servers reserved exclusively for hosting private cloud workloads. All servers were based on blade servers with necessary components.
- **Storage Cluster:** External all-flash storage arrays and storage controllers were evaluated. The storage area network (SAN) was based on dedicated network components.

Figure 4: Traditional 3-T

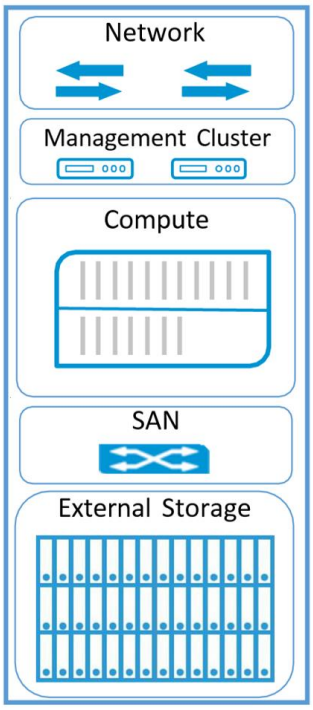

ier

The traditional 3-tier based system we evaluated supported a virtualized private cloud on top of the vendor's pre-built and/or reference architecture systems. Most vendors ship VMware capable systems with pre-installed vSphere and vCenter components. The creation of an SDDC-like environment may require either a professional services engagement or significant customer-led customization and will result in limited integration, automation, and lifecycle management capabilities.

TRADITIONAL 3-TIER ARCHITECTURE INFRASTRUCTURE MANAGEMENT

An additional overhead of 3-tier systems is that each system component comes with its own management software. This makes sense given that each device stands on its own; however, it leads to a complex set of individual tools with varying degrees of overlap and unique capabilities. For instance, most vendors support vCenter plug-ins but typically not all functionality and troubleshooting is available through these plug-ins. This adds complexity and the customer must be fully trained on and understand each tool. Several 3-tier companies have been trying to mitigate this by creating their own uber-tool to manage each of the 3 tiers as if they were one. While this is a good first step, these types of tools have yet to eliminate the need for individual tools. The bottom line: traditional 3-tier based architectures require a more complex system management toolset.

Summary of Key Components for the two Private Cloud Alternatives*

	VMware Cloud Foundation *(Basic Edition)*	**Traditional 3-Tier DIY or CI**
SOFTWARE		
Compute Virtualization	*VMware vSphere Ent+*	*VMware vSphere Ent+*
Network Virtualization or Software Defined Networking	*VMware NSX DC PRO*	*VMware NSX DC PRO*
Storage Virtualization	*VMware vSAN ADV*	*None*

PO Box 117 - West Dennis, MA 02670 -T: 508.398.0075 - F: 508.398.0071

SOFTWARE for MANAGEMENT		
Virtualization Management	*vCenter*	*vCenter*
Lifecycle Management	*VMware Cloud Foundation*	*Software products from multiple vendors*
Infrastructure Management & Automation	*VMware Cloud Foundation*	*Individual element manager plus vCenter plug-ins*
HARDWARE		
Network	*Top of Rack Switches* *Management switch*	*Top of Rack Switches* *Management switch*
Security and Load Balancing	*None* *(all within NSX)*	*None* *(all within NSX)*
Storage Area Network	*None* *(all within vSAN)*	*Independent FC SAN switches*
Compute for workloads	*Modular scale-out HCI servers* *(Any Qualified vSAN ReadyNode)*	*Blade servers* *Blade chassis* *Additional components*
Compute for Management cluster	*Modular scale-out HCI servers*	*Management cluster built with standard rackmount servers*
Storage and Controllers	*Modular scale-out HCI servers and within vSAN*	*External all-flash arrays and storage controllers*

**To completely build out a fully functional private and hybrid cloud, we recommend adding to these scenarios by layering on cloud automation software such as VMware vRealize Suite. The focus of this study was to compare just the essential core technology differences between scenarios.*

COST COMPARISON METHODOLOGY AND APPROACH

For the cost comparison analysis, the systems were configured and optimized for workload use cases where the quantitative hardware metrics could be calculated and compared to reference architectures from the vendors. As a result, each 3-tier system was tuned to meet similar CPU, memory, and storage criteria such that comparative study was possible while staying within the reference architecture guidelines. This allowed us to compare usable VM capacity fairly, not handicapping a product configuration with too little memory or storage. The following are the guiding principles used in this comparative research:

- Optimize system sizing for an infrastructure deployment of around 1,000 moderately sized VMs.
- Use a vendor's reference architecture (where available) for sizing and best practices.
- Optimize each solution with similar compute, storage and networking resources to make the comparison as fair and balanced as possible.
- Focus on the key CapEx components and associated 3 years of ongoing support costs. Ignore cost components that show little to no variation between the architectural approaches (e.g. no significant difference in power consumption or floor space utilization was found).
- Taneja Group validated a VMware OpEx savings study and model. VMware conducted the study in its own data center and in VMware Cloud Foundation customer datacenters by comparing labor and time savings versus the traditional 3-tier systems that the customers were using.

Assumptions and Configurations for each Scenario

As discussed above, each traditional 3-tier scenario was configured to be as close as possible to the VMware Cloud Foundation scenario. Further, since none of the traditional 3-tier systems had a fully integrated SDDC software stack, Taneja Group added the missing SDDC software, making it as close as possible to the VMware Cloud Foundation integrated software stack. Since a common load testing tool is not widely recognized for effective VM density, the approach to calculating the total usable VM

capacity was first to normalize all the available CPU capacity dedicated to guest VMs, and then give each solution the same total usable VM capacity based on a consistent number of vCPUs per VM. The goal was to have each comparative solution approximate the same equivalent usable VM capacity per system. The following table gives the high-level configurations for each scenario:

System Configuration Comparison Matrix		
Component	**VMware Cloud Foundation**	**Traditional 3-Tier DIY or CI**
Cloud Platform	VMware Cloud Foundation *(Basic Edition)*	vSphere and additional SDDC software
Compute Platform	HCI Rackmount Servers (all nodes the same)	Blade Servers (8/16 nodes per chassis)
Cores per CPU	18	18
Memory per Node	512 GB	512 GB
Dedicated Management Nodes in Solution	4 (from HCI servers above)	4 (isolated rackmount servers)
Available Compute Nodes in Solution	28	28
Storage Technology	VMware vSAN	External all-flash array, SAN switches, and storage controllers
Raw SSD Storage Capacity	269 TB	184 TB
Raw-to-Effective Capacity Overhead Estimate	50%	80%
Effective Capacity	134 TB	147 TB

Assuming the above configurations, the usable VM characteristics were calculated based on the following formula as outlined in the table below:

VM Capacity Calculation		
Item	**Traditional 3-Tier**	**VMware Cloud Foundation**
Total CPU Cores in Solution	1,152	1,152
Dedicated Cores for Management	144	144
Compute Cores Available for Customer VMs	1,008	1,008
Overcommit Ratio	4.0	4.0
Available vCPUs	4,032	4,032
VMs Supported (based on 4 vCPUs per VM)	**1,008**	**1,008**
Estimated Memory Available per VM (GB)	14.2	14.2
Estimated Usable Storage per VM (GB)	146	133

HARDWARE AND SOFTWARE COST ANALYSIS

The approach used to calculate the cost metrics for each of these scenarios was to build solutions that matched the characteristics above. Where possible, for the 3-tier systems, the HW and SW BOM structures exactly matched those specified for each vendor's reference architecture. The list prices were obtained from vendor pricing tools, vendor reseller quotes, and pricing databases. Because of the constraints on building real systems based on real pricing, not every dimension in the area of

compute/storage/memory ratios could be matched exactly. Instead of trying to force-fit the ratio, we leave it up to the reader to gauge the importance of each variable in comparing the scenarios. The costs of the solutions were calculated in the following categories:

Hardware CapEx

- For compute/server/networking, the actual prices of components were used based on current vendor pricing and vendor reseller quotes.
- All hardware-specific management packages (those provided by the hardware vendor or the software bundled with the hardware specifically to run on that appliance) were included in the hardware CapEx section. In many cases, the software was bundled with an appliance and could not be separated – these were included with the hardware.

Hardware OpEx

- The actual maintenance pricing was used where available, and if a three-year quote was not available, the maintenance was extended on a per year basis. If a particular hardware component maintenance quote could not be found, a 5%-8% per year hardware maintenance cost was assumed, which fit within the range of the actual maintenance quotes received.
- In a previous cost comparison analysis, Taneja Group validated an OpEx study contrasting labor costs between a VMware Cloud Foundation based solution and a Traditional 3-Tier solution. From that study, we concluded that a VMware Cloud Foundation-based solution would consume .5 Full Time Equivalent (FTE) of labor and a Traditional 3-Tier approach would consume 1.5 FTE to maintain just the hardware portion of the solution.

Software CapEx

- All perpetual software components that were not specifically tied to the hardware vendor were included based on the research methods described above.
- Miscellaneous software components that were deemed similar across all scenarios were not included. For example, the costs of Microsoft Windows software for the management cluster or client VM hosted software were not included.

Software Support

- Annual software support was included and calculated for three years.

Incremental On-premises Labor

- Validated from a previous study for the relative difference in labor cost to manage the infrastructure portion of traditional 3-tier versus VMware Cloud Foundation SDDC.

Cost Comparison Results – Upfront Costs for HW, SW, and Support

The table below summarizes the upfront cost attributes of each system for a private cloud.

PO Box 117 - West Dennis, MA 02670 -T: 508.398.0075 - F: 508.398.0071

TCO Component Costs for HW, SW, Support (public prices) For a 1008 VM Deployment Scenario		
Cost Item	**Traditional 3-Tier**	**VMware Cloud Foundation**
HW CapEx	$2,668,202	$1,586,615
Appliance/Compute	*$1,134,711*	*$1,394,846*
Storage and SAN	*$1,336,090*	
Network	*$130,768*	*$123,945*
Rack and Cables	*$27,804*	*$36,038*
Additional SW	*$2,796*	*$0*
Installation	*$36,033*	*$31,786*
SW CapEx	$549,310	$853,310
HW Maintenance (3 Years)	$813,053	$334,140
SW Maintenance (3 Years)	$411,983	$639,983
Incremental On-premises Labor	$585,000	$195,000
Total Costs over Three Years	**$5,027,547**	**$3,609,048**

Our in-depth analysis indicates that using VMware Cloud Foundation results in significant upfront cost savings over using traditional 3-tier systems. The combined costs of the server and shared storage were significantly reduced by using hyper-converged infrastructure, similar to what can be achieved in large web-scale service provider datacenters. The resulting hardware cost savings help to offset the incremental SDDC software costs needed to provide storage and networking capability that leading 3-tier component vendors typically provide within hardware.

Upgrading to an enhanced Hybrid Cloud

The private clouds outlined above can be enhanced by layering VMware's multi-cloud automation software called VMware vRealize Suite on top of each system. The VMware vRealize Suite is included in all but the Basic version (i.e. Standard, Advanced and Enterprise versions) of VMware Cloud Foundation, providing increasing levels of hybrid cloud automation capabilities. These upgraded versions also deliver steadily more advanced vRealize Suite functionality along with enhancements in networking (NSX) and storage (vSAN) capabilities.

Similarly, these enhancements can also be layered on top of the traditional 3-tier private cloud by individually adding the same versions of vRealize Suite to that system and upgrading the NSX version to match. It was outside the scope of this paper to go in-depth on nuances of the various multi-cloud or hybrid cloud capabilities that each of these enhancements brings.

THE CONTRASTING APPROACHES TO MOVING TO A PUBLIC CLOUD

In this section of our TCO analysis we are focusing on moving the same 1,008-VM workload we had analyzed for private clouds to two public cloud alternatives. The approach we took in this scenario was to build public cloud configurations that could handle a set of IaaS workloads similar to those used in the two previous private cloud alternatives. Our research shows that for most customers the goal will be to run in a hybrid cloud configuration by connecting their private cloud to a public cloud.

The first public cloud we analyzed was VMware Cloud on AWS. VMware Cloud on AWS has been built on VMware's Cloud Foundation core software and can be deployed as-a-service on AWS as easily as one can do a simple mouse click. The difference now is that the hundreds of thousands of VMware customers that have come to rely on VMware as their key enterprise virtualization provider can instantly get a fully functional hybrid cloud with all the security, control, and features they depend on in their on-premises VMware environments. Also, customers will enjoy seamless workload migration

PO Box 117 - West Dennis, MA 02670 -T: 508.398.0075 - F: 508.398.0071

from private to public clouds, advanced disaster recovery capability, and—by being on AWS public cloud—safe and secure access to additional AWS services.

The second public cloud we analyzed was Microsoft Azure. Microsoft Azure is a popular public cloud infrastructure and is the most frequently cited public cloud alternative to AWS. When moving on-premises workloads to a public cloud that is not built on VMware, we found that there is a significant refactoring cost to either migrating (lifting and shifting) or reengineering applications to run on Azure. To quantify these incremental costs, we performed in-depth primary research to gauge the cost of refactoring applications to the hyper-scaled public clouds.

VMware Cloud on AWS architecture explained

Unlike most public cloud infrastructure, VMware Cloud is transparent. VMware Cloud on AWS uses a hyper-converged architectural approach, powered by VMware Cloud Foundation, which is used by many of the leading HCI and private cloud deployments on premises today. When deployed on AWS, the on-demand VMware Cloud service enables you to run applications across vSphere-based environments and includes access to a broad range of AWS services. The service integrates vSphere, vSAN and NSX along with VMware vCenter management, and is optimized to run on dedicated, elastic, bare-metal AWS infrastructure. Businesses can manage their dedicated public-cloud resources with familiar VMware tools. Figure 5 is a graphical representation of VMware Cloud on AWS architecture.

Figure 5: VMware Cloud on AWS high level architecture

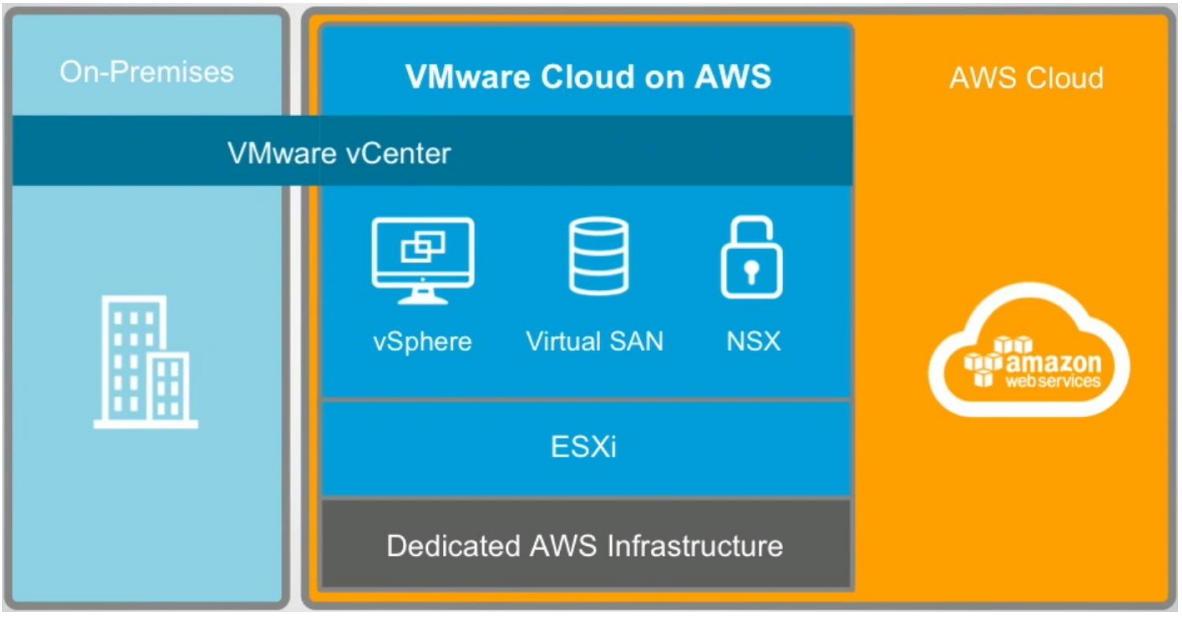

SOURCE: VMWARE

The baseline VMware Cloud on AWS service comes with four dedicated compute nodes configured in a highly resilient hyper-converged cluster. The service can scale on-demand one node at a time to a maximum of 16 nodes for a single HCI cluster environment. Businesses needing additional scale can then stand up a new cluster starting again with a minimum of three nodes. The following table outlines exactly what is included from a hardware infrastructure perspective.

VMware Cloud on AWS	
Component	**Description**
Cloud Software Platform	VMware Cloud Foundation Core Software
Cloud Compute Platform	AWS bare-metal Servers (all nodes the same)
Cores per Node	36
Memory per Node	512 GB

PO Box 117 - West Dennis, MA 02670 -T: 508.398.0075 - F: 508.398.0071

NVMe-based SSD Raw Capacity per Node	10.7 TB
Storage Technology	VMware vSAN
Minimum Nodes in a Cluster	3
Maximum Nodes in a Cluster	16
Networking Technology	VMware NSX

Following are some of the key attributes that make VMware Cloud on AWS an ideal hybrid cloud solution.

- **Operational consistency across on-premises and public cloud.** Familiar VMware tools such as vCenter, vSphere API, and vSphere clients make managing VMware infrastructure on AWS exactly like you do on-premises. No costly workforce training is needed to get started on a hybrid cloud journey. Hybrid Linked Mode allows you to link public and private vCenter server instances together and thus manage the inventories of both environments from a single vSphere Client interface.
- **Integrated hybrid cloud management.** Simplified orchestration and operation management for a hybrid cloud environment is provided using the full vRealize suite of applications from VMware. Whether you use VMware-based hybrid cloud tools or a myriad of other third-party tools, you'll benefit from the proven interoperability these tools have demonstrated with vSphere-based virtualization environments.
- **Enterprise-grade application workload support.** VMware Cloud on AWS leverages enterprise-class, bare-metal compute nodes that support the latest modern processor complex architectures, combined with locally attached SSD technology that supports the NVMe protocol. VMware vSAN software-defined storage combined with NVMe-based SSDs creates a shared storage environment with a long track record of supporting mission-critical workloads. Having an all-flash storage infrastructure means businesses will not have to guess what workloads may or may not be suitable to move to VMware Cloud on AWS.
- **Higher-than-average VM density and efficiency.** VMware vSphere environments are well known for their resource efficiency. You can take full advantage of this efficiency when deploying workloads on VMware Cloud on AWS. Businesses do not know how much public cloud vendors overprovision resources or how efficiently their underlying virtualization technology performs. With VMware Cloud on AWS transparent approach, you can pack more VM instances in the environment without paying any incremental costs, resulting in a more efficient and cost-effective solution. Unlike other public cloud approaches, increased density and efficiency leads to a lower cost per VM, versus just boosting the profitability of the cloud provider.
- **Seamless network bridging and features.** By leveraging NSX technology across public and private cloud environments, VMware Cloud on AWS solves many complex compatibility issues for hybrid clouds. VMware HCX makes it possible to use vMotion to move workloads across clouds, just as could previously only be done across server nodes.
- **Uncompromised workload portability.** Arguably one of the most over-hyped and under-delivered promises of hybrid clouds is the capability to burst workloads seamlessly to a public cloud. With VMware Cloud on AWS, this capability is delivered by leveraging some of the most cherished VMware features that now work seamlessly between VMware environments on-premises and in the public cloud. Such features include vMotion to transparently move a workload and DRS to automatically load-balance between VMware Cloud on AWS and on-premises vSphere environments.
- **Single-contact premium support included.** Support in public clouds can vary depending on the level of service you pay for. With VMware Cloud on AWS, VMware manages and operates

the service and provides full lifecycle support of all the components in the solution (e.g. emergency patches, software upgrades, etc.). Customers will benefit from world-class, 24X7 support services including an online support center with FAQs, forums & chat capability.

- **Direct Access to AWS Services.** Applications running in VMware Cloud on AWS have the same high-performance (low latency, high bandwidth) access to AWS services (e.g. S3, AI, analytics, Lambda, etc.) as if those applications were running out of native EC2. This is due to networking integration between VMware and AWS infrastructures that provides direct dedicated access to AWS environment and services.

VMware Cloud on AWS efficiency and pricing expectations

VMware enables customers to build the right strategy to take advantage of hybrid cloud. Everything is transparent. Customers can now leverage VMware capabilities they've successfully used on-premises, such as VM density/oversubscription, vMotion, and HA, and apply these capabilities in a hybrid cloud architecture. VMware is changing the metrics of cloud pricing with this offering. While actual efficiency calculations depend on the overall number of workloads deployed in the cloud, Taneja Group found that the following table is a good rule of thumb for realistic efficiency expectations when using VMware Cloud on AWS.

VMware Cloud on AWS Efficiency Calculation		
Component	**Value**	**Description**
Cores per Node	36	Physical cores available per computer node
Overcommit Ratio	4	Very conservative approach when considering effects of Hyperthreading.
Available vCPUs per Node with oversubscription	144	vCPUs are the standard unit of compute resource commonly ordered in public clouds.
Memory available per vCPU	3.5 GB	This is a generous amount of memory per vCPU for the workloads.
Available effective SSD Capacity per vCPU	47 GB	Based on a RAID 6 configuration and a conservative 1.44:1 deduplication ratio.

VMware provides multiple pricing options when provisioning VMware Cloud on AWS. You can rent infrastructure on-demand, or for better pricing, reserve infrastructure in one-year or three-year increments, depending on your needs. Discounts go up as the reservation period increases. The following table outlines pricing at the time this report was published. Pricing includes a robust bundle of enterprise-level service and support.

VMware Cloud on AWS Pricing Calculation			
Pricing Component	**On-Demand (Hourly)**	**1 Year Reserved**	**3 Year Reserved**
List Price ($ per host)*	$8.3681/hour	$51,987/year	$109,366/3 year
Effective Hourly**	$8.3681/hour	$5.9346/hour	$4.1616/hour
Savings Over On-Demand		30%	50%
**Pricing based on US West Data Center.* ***Normalized to hourly for comparison only.*			

VMware Cloud Specific Configuration and Costs

The configuration sizing for VMware Cloud on AWS is very straightforward. Rather than going through a pricing calculator with hundreds of choices and combinations, we only had to calculate how much VM capacity each host node could handle and then multiply the number of hosts needed to cover our workload scenario. This node-based pricing concept is very important because unlike the Azure scenario, we can change VM configurations on the fly and still know exactly what the monthly cost envelope will be. The follow table outlines the configuration and cost for VMware Cloud on AWS.

PO Box 117 - West Dennis, MA 02670 -T: 508.398.0075 - F: 508.398.0071

VMware Cloud on AWS Public Cloud Pricing Calculation		
Cloud Component	**Selections**	**3-Year Reserve Cost**
VMware Cloud Host	36 Cores, 512 GB Memory, 10.7 TB raw SSD Storage	$3,038/Month
IaaS Workload Scenario	28 Hosts (1,008 VMs)	$15,190/Month
Support	Premium	Included
Total 3-Year Cost		**$3,062,248**
Cost per VM Per Month		**$84.39**

Microsoft Azure Public Cloud Alternative

We looked at Microsoft Azure as the second public cloud in our cost comparison analysis. The cloud scenario was configured as general-purpose IaaS systems capable of running business-critical workloads (e.g., with high availability and 24x7x365 service response times). We relied on Microsoft recommendations for VM sizing and storage configuration and included extended service contracts that could hit this level of SLA. Since Microsoft provides many different VM sizes, we first picked an Azure VM size that met our criteria and then matched that VM size on VMware Cloud on AWS. We used this approach to make each public cloud environment in our analysis as close to equivalent as possible. The following is a summary of our comparison approach and assumptions.

- Optimize IaaS sizing, making the environment large enough to support the shift of a reasonable set of business-critical workloads to the public cloud. Start with an Azure-capable VM instance and then match to VMware Cloud.
- Do not include costs that should be relatively equal between the two scenarios (e.g., AWS versus Azure egress fees)
- No assumption around what it will take to get the workload into the public cloud. This part of the analysis will be examined in the refactoring cost analysis.
- Taneja Group research indicates that businesses that plan to shift their critical workloads to the public cloud intend this as a long-term solution. Therefore, the length of the analysis is based on a three-year lifecycle.

Configurations for each Public Cloud

The following table shows the relative configuration size and services purchased for each public cloud.

VMware Cloud on AWS Pricing Calculation			
Cloud Component	**Minimum Required**	**Azure Selections**	**VMware Cloud Selections**
Total VM Instances	1,000	1,008	1,008
vCPUs Per VM	4	4	4
Memory per VM	8 GB	8 GB	14 GB*
Guest O/S	Windows	Windows	Windows
Average Storage per VM (SSD-Based)	125-150 GB	128 GB	148 GB
**VMware Cloud on AWS memory was greater than the minimum required as vCPUs were the limiting factor in the configuration. You could easily configure a workload that needed more storage and fewer vCPUs on VMware Cloud.*			

PO Box 117 - West Dennis, MA 02670 -T: 508.398.0075 - F: 508.398.0071

Azure Specific Configuration and Costs

With any public cloud, there is a myriad of vCPU size and memory configuration options to choose from. For the IaaS scenario on Azure, we used the Microsoft public pricing calculator and went through the following selection process:

1) Select a Virtual Machine type and size
2) Select durable, permanent storage that will meet the workload requirement
3) Select additional support services

This process yielded the following selections and pricing for our scenario:

Azure Public Cloud Pricing Calculation		
Cloud Component	**Selections**	**3-Year Reserve Cost**
Virtual Machine Type	F4: 4 vCPU(s) 8 GB RAM 64GB Temp Storage	$68.86/Month
Permanent Storage	Premium Managed 128 GB SSD	$17.92/Month
Snapshot Storage	12 GB	$1.44/Month
IaaS Workload Scenario	1,008 VMs	$15,662/Month
Support Costs	Professional Direct*	$1,000/Month
Total 3-Year Cost		**$3,273,327**
Cost per VM Per Month		**$90.20**
**Professional Direct is based on $1,000 per month per customer regardless of number of VMs supported.*		

While this three-step process appears rather simple, we found that numerous options were available to us in each step, requiring significant time to determine which option made the most sense for our specific workload. Any customer not familiar with Azure will similarly need to invest considerable time upfront to research and sort out the various options.

Refactoring Applications: The hidden cost of moving to hyper-scale public clouds

One of the biggest barriers to corporate adoption of public clouds has been the time, effort and cost to move and deploy applications there. Apart from applications developed to run natively in the cloud, each new application destined for the cloud must at least be migrated and adapted, and in many cases also reengineered, to work successfully in each public cloud environment. Taneja Group research suggests that companies often underestimate the migration and refactoring investment required to get their business applications up and running. In a 2018 survey, we asked a series of questions to try to determine and quantify the hidden costs of refactoring, based on the types of applications being moved and the estimated effort and time invested.

In our research we surveyed over 200 IT practitioners who have migrated on-premises workloads to native AWS or Azure public clouds. We knew it would be a challenge a come up with a quantifiable number to assign to the cost of migrating workloads from an on-premises environment to a hyper-scaled public cloud. We therefore decided to ask these practitioners to respond to costs using three different methods to see if we could see some convergence on the costs. We also asked them to differentiate the costs for migrating workloads using packaged applications that need little or no engineering changes – commonly called "lift and shift" or "replatforming" – from those custom applications that require some level of engineering changes, which we called "refactoring". What we found was a remarkable consistency in the cost-per-VM to migrate workloads to the cloud. Following are the 4 key questions we asked respondents to gather the information needed to complete our analysis.

1. How many VMs were migrated to the public cloud?
2. What was the estimated cost to move these VMs to the public cloud?
3. What was the estimated people time to move these VMs to the public cloud?
4. What was your approximate cost per VM to move these VMs to the public cloud?

In our analysis we found that while some very large and expensive projects skewed the cost per VM toward the high end, some other respondents claimed that their migration cost nothing. We took the approach of throwing out both these high and low instances to come up with a normalized average for the costs. The following table summarizes the costs using each of the methodologies.

Cost Analysis: Migrating Workloads to the Public Cloud			
Item	**Migrate Packaged Apps "Replatforming"**	**Migrate Custom Apps "Refactoring"**	**Average Cost**
Number of VMs migrated to Cloud	802 (average)	749 (average)	NA
Cost Method One: Total Project Cost divided by number of VMs migrated	$1,021/VM	$958/VM	$990/VM
Cost Method Two: People Cost* divided by number of VMs migrated	$911/VM	$968/VM	$940/VM
Cost Method Three: Cost per VM as stated by the respondent.	$642/VM	$1,337/VM	$989/VM
**For People Cost we used $120K per year average.*			

We were encouraged that each of these three calculation approaches converged on similar numbers for migration costs. For the sake of this paper we felt the most accurate cost estimate was based on method three, so we used the $989/VM in our final comparative analysis.

Comparing VMware Cloud to Azure Public Cloud Offerings

The follow table and graph outline the comparison of VMware Cloud on AWS versus a comparable Azure configuration that can handle a similar workload and SLA expectations for 1,008 VMs. Our assessment includes the one-time cost of refactoring your applications to run on the public cloud, which you can expect to incur in moving applications to Azure or other non-VMware Cloud Foundation based public clouds.

Azure versus VMware Cloud Public Cloud Pricing			
Item	**Azure**	**VMware Cloud on AWS**	**VMware Cloud Savings**
3-Year Cloud Subscription Cost	$3,273,327	$3,062,248	$211,079 (6%)
One-time Migration/Refactoring Cost	$996,912		$996,912 (100%)
Total Cost	$4,270,239	$3,062,248	$1,207,991 (28%)

PO Box 117 - West Dennis, MA 02670 -T: 508.398.0075 - F: 508.398.0071

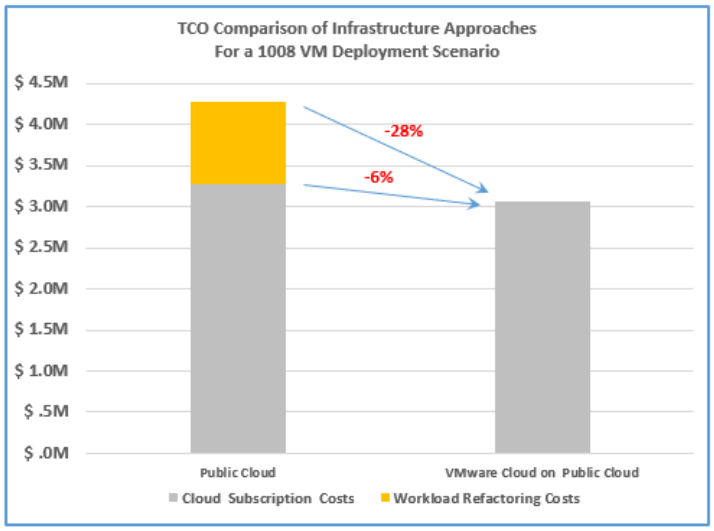

CONTRASTING THE FOUR CLOUD ALTERNATIVES

Whether deployed on-premises as a private cloud or via a publicly available provider service, cloud IaaS is rapidly becoming the infrastructure of choice as customers increasingly look for ways to save costs. The table and graph below present in more detail the comparative costs of the four cloud alternatives over a period of three years, each capable of handling 1,000 VMs running similar types of workloads.

Cost Item	Traditional 3-Tier	VMware Cloud Foundation	Public Cloud	VMware Cloud on Public Cloud
On-Premises Up Front Costs	$3,217,512	$2,439,925		
On-Premises Ongoing Costs	$1,810,035	$1,169,122		
Cloud Subscription Costs			$3,273,327	$3,062,248
Workload Refactoring Costs			$996,912	
Total Costs over Three Years	**$5,027,547**	**$3,609,048**	**$4,270,239**	**$3,062,248**
Assumed VM Density	1,008	1,008	1,008	1,008
Cost per VM	**$4,988**	**$3,580**	**$4,236**	**$3,038**
Savings versus Traditional 3-Tier Approach		**28%**	**15%**	**39%**

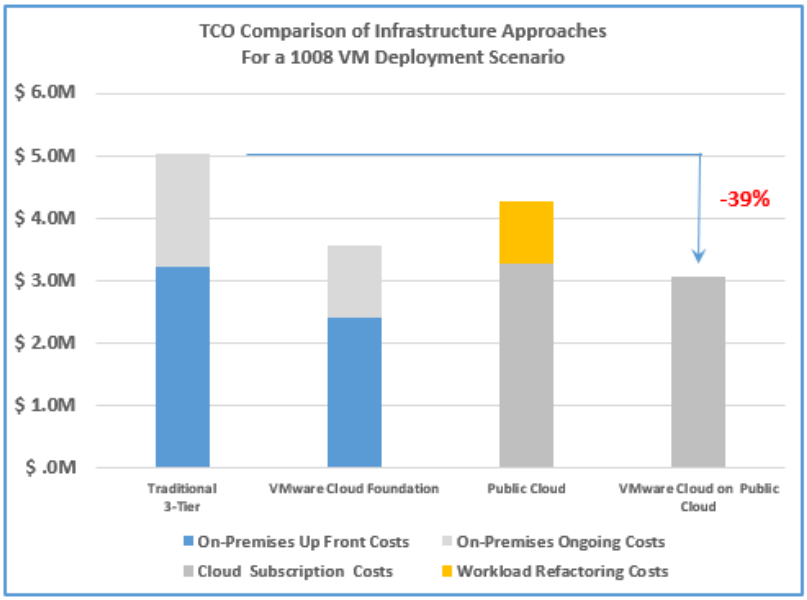

PO Box 117 - West Dennis, MA 02670 -T: 508.398.0075 - F: 508.398.0071

Each cloud scenario we analyzed comes with tradeoffs between costs and benefits. Here is a summary of qualitative factors that may lead you to choose one cloud approach over another.

Reasons to Choose Different Cloud Types

Choosing the Right Cloud Approach		
Cloud Type	**Reasons Customer Choose**	**Some Disadvantages**
Traditional 3-Tier Private Cloud	• Prefer Blade Server Technology for compute • Storage versus Compute scaling occurs independently • Need to support bare-metal workloads using the same infrastructure • Need a very specific optimized storage performance characteristic	• Can cost much more in upfront CapEx investment • Typically, the infrastructure is heavily overprovisioned to start • Upgrade cycle tends to be lumpy and comes in 3-5 year increments
VMware Cloud Foundation Based Private Cloud	• Ability to match private cloud to public cloud architectures • Ability to buy only what you need upfront and add as you need more • Lowest cost for on-premises infrastructure	• Some types of workloads may have specific compute or storage needs that may or may not be well supported • Compute versus Storage sizing and capabilities not always perfectly optimized
Public Cloud	• Have large workloads that ebb and flow up and down (e.g. due to seasonal variations) • Good for provisioning globally accessible, consumer-facing web-apps or content • Need access to other public cloud services (AI cloud services) • Startup companies tend to start first in the public cloud	• Heavy refactoring or replatforming costs to move on-premises workloads to public cloud • Can cost much more if data needs to move out of the cloud • Specific applications may need to be on-premises due to regulatory compliance demands
VMware Cloud on Public Cloud	• Want a seamless hybrid cloud environment with on-premises VMware technology • Lowest cost public cloud for stable business-critical workloads • Virtually no replatforming costs needed to move on-premises workloads to public cloud • Easiest way to transparently migrate workloads back and forth between public and private clouds	• Does not scale in as fine-grain increments as public cloud

Whether driven by TCO or these other factors, we believe each of these four private and public cloud based approaches will continue to be adopted well into the future.

TANEJA GROUP OPINION

Hybrid clouds are no longer just a wave of the future. They are here today and here to stay. As we learned in this research study, a majority of enterprise decision makers favor hybrid cloud as their long-term architecture, and many companies are already well down the path of deploying them.

When looking at hybrid cloud options, we think it is best to focus on solutions that maximize your choice of private and public cloud alternatives, to minimize the potential for provider lock-in and ensure you can deploy each of your workloads in the cloud that provides the best fit of features, function and cost. Hybrid solutions that offer full cross-cloud compatibility and workload portability further boost your choice of deployment and enable workloads to run in true hybrid fashion.

PO Box 117 - West Dennis, MA 02670 -T: 508.398.0075 - F: 508.398.0071

In several past studies, most recently in late 2017, Taneja Group validated the advantages of software-defined data center architectures over traditional 3-tier hardware-driven deployments, and in particular the cost and agility advantages of VMware Cloud Foundation-based clouds over other public and private cloud alternatives, such as those based on Microsoft Azure and Azure Stack. And as we've seen in this latest study, when app migration and refactoring costs are considered, the relative TCO advantage of VMware Cloud Foundation-based clouds grows even further.

As you consider different hybrid cloud options, we recommend you evaluate solutions based on VMware Cloud Foundation technology, including a VMware Cloud Foundation-based private cloud and VMware Cloud Foundation-enabled public clouds such as VMware Cloud on AWS or other VMware Cloud solutions supported by VMware partners. These clouds will lay the groundwork for a fully capable and cost-effective hybrid cloud deployment that you can continue to build on in the years to come.

Taneja Group Note: This paper contains ideas on calculating real-world TCO for the presented solutions. However, differences in customer environments, skill sets, and related services will drive differences in TCO analyses. Further, the supported business outcomes and especially the automation level of the compared offerings vary. You will need to conduct an analysis tailored to your own environment to develop TCO numbers consistent with your environment. Finally, the analysis presented in this paper is designed to raise awareness about how to evaluate solutions. The analysis is not about proving which solution is the best for specific circumstances. Each solution covered in this analysis has its own proven merits.

PO Box 117 - West Dennis, MA 02670 -T: 508.398.0075 - F: 508.398.0071