

White Paper

# Understanding the True Cost of Public and On-premises Cloud Solutions

## The Economics of Cloud-based IT Agility

By Terri McClure, ESG Senior Analyst, Brian Garret, Vice President, Validation Services, and Leah Matuson, Research Analyst

January 2017

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

## Contents

Introduction .....	3
How Do You Create Cost-effective, Agile IT? .....	4
Influential Factors .....	4
Issues to Address when Building a Cloud Environment.....	4
Private Cloud Infrastructure Approaches .....	5
Do-it-yourself (DIY) 3-tier Infrastructure .....	5
Converged Architecture .....	5
Hyperconverged Architecture.....	5
Buying On-premises Infrastructure versus Renting in the Cloud.....	6
Fundamental Questions for a True Cost-to-own Discussion .....	7
On-premises Solutions .....	7
Public Cloud-based Solutions.....	7
Comparing the Cost of Public Cloud and Private Cloud Powered by Intel and Lenovo.....	8
The Bigger Truth.....	11

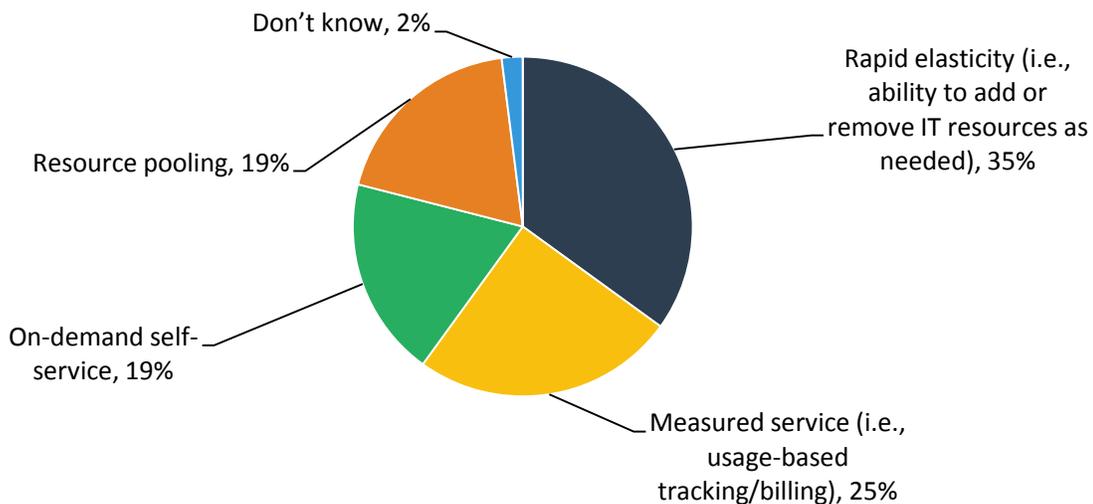
## Introduction

Businesses across the board are recognizing that increasing IT agility can have a direct impact to the business. Achieving IT agility enables the organization to quickly respond to today’s rapidly evolving business and competitive environment. So it stands to reason that many enterprises are closely examining strategies to drive or increase IT agility. But how do they achieve it? Many organizations are turning to the public cloud.

The public cloud movement is often being driven from the top down based on a fundamental belief that the only way to attain agility while driving down costs is to move to the public cloud and rent resources as needed. But more frequently, IT organizations are realizing that tremendous IT agility and cost savings can be gained from building a private cloud in house. In fact, more than a third of the IT managers surveyed by ESG about key private cloud capabilities cite elasticity of resources as the most important capability.<sup>1</sup> This elasticity of resources is often what IT organizations are looking for from the public cloud, but advances in technology such as converged and hyperconverged systems can help provide the elasticity IT needs in a private cloud. There is more to think about, though, when it comes to the decision to deploy a public or private cloud.

**Figure 1. Most Important Private Cloud Capabilities**

**Which of the following do you believe is the most important capability of a private cloud infrastructure? (Percent of respondents, N=308)**



Source: Enterprise Strategy Group, 2017

The public cloud is certainly an attractive option. Deploying applications in the public cloud eliminates the physical administration and ownership of physical resources, and IT can plug in to a browser-based, services-oriented experience using APIs. So public cloud can offer agility and savings. And certainly initial entry costs are attractive for public cloud services—no large upfront CapEx investment is required like there is when you build in house. However, entry cost is only part of the story: IT professionals comparing public and private cloud must also consider factors such as risk, cost of scale, and cost over time.

Moving to the cloud does introduce some level of risk. The organization becomes dependent on the service provider to secure and protect data. And while some service providers are transparent about how they maintain security, others are

<sup>1</sup> Source: ESG Research Report, [The Cloud Computing Spectrum, from Private to Hybrid](#), March 2016.

not. Also, the regulatory environment needs to be considered—knowing (and specifying) where data is stored, who can access it, how long it must be retained, and whether it is dispositioned properly is key. This is true for both primary and backup data. Many countries specify that certain types of data cannot be stored outside their geographical boundaries. This can make choosing a cloud provider onerous for organizations with business in multiple geographic regions, and could lead to the need to subscribe to multiple cloud providers, adding administrative overhead. And still, lack of visibility into the entire data environment makes it hard for IT to parse what is “safe” for the cloud and what should be kept on-premises to ensure regulatory requirements are met. So although businesses need agility, it must be balanced with risk.

Cost of scale and cost over time also needs to be considered. When looking at costs associated with scaling an application, organizations need to dig deep into cloud service providers’ pricing models to understand how adding resources, such as virtual machines, cores, memory, or storage, will impact cloud costs over time. These costs can add up quickly and are often not sufficiently considered when making an upfront cloud investment. IT organizations need to do the math and figure out what it really costs to subscribe to a cloud provider for the application lifecycle versus building a private cloud. Considering time and scale, building a private cloud can be more cost effective than subscribing to a public cloud service.

So how can organizations get the performance, security, and agility of a public cloud *without* the risk, and what does it really cost? Lenovo is engineering solutions, based on the industry’s most reliable Intel Xeon servers, that can deliver a public cloud experience in a private cloud environment. Understanding the cost equations and where the Lenovo solutions make sense is the focus of this paper.

## How Do You Create Cost-effective, Agile IT?

At the core of this discussion are two primary choices: outsource to the public cloud (rent resources), or build in house (create an on-premises cloud). There is a strong *perception* that outsourcing to the public cloud can create significant savings. In fact, according to ESG research, outsourcing to the public cloud is the most-cited cost savings initiative among enterprises surveyed.<sup>2</sup> But the cloud may not be as cost effective as most organizations believe. That’s why it is important to understand the broader long-term implications of your public versus private cloud decision.

Outsourcing to the public cloud is fast, but IT is still on the hook for managing the environment, developing and maintaining applications, integrating the system into the broader IT security and application ecosystem, and sustaining relationships. Outsourcing to the public cloud still requires monitoring the workload, and ensuring applications perform as expected.

### Influential Factors

Technology is constantly evolving, and those technology capabilities originally associated with cloud service providers are now broadly available. With technology innovations such as virtualization, and converged and hyperconverged infrastructure, IT can now create a highly agile on-premises cloud infrastructure that delivers IT-as-a-service to the organization.

## Issues to Address when Building a Cloud Environment

You’ll need to address several basic issues when building cloud environment, including resource ownership, management, and location. This means you have some important decisions to make:

- **Who owns it?** You can own your infrastructure, or the cloud provider can own it.
- **Who sources and builds it?** You can source and buy the parts, or the cloud provider can source the parts.

---

<sup>2</sup> Source: ESG Research Report, [2016 IT Spending Intentions Survey](#), February 2016.

- **Who manages it?** No matter who owns it, you can manage it, or you can have someone else manage it.
- **Where is it housed?** It can live on-premises, or at a cloud provider.

## Private Cloud Infrastructure Approaches

Should you choose to deploy a private cloud infrastructure on-premises, and buy, source, manage, and house the cloud environment, three approaches can be used to build the cloud foundation: do-it-yourself, converged, and hyperconverged. These approaches are not mutually exclusive—indeed, many companies are using a combination of the three. Understanding the differences between the three, and how and where to deploy each is key.

### Do-it-yourself (DIY) 3-tier Infrastructure

Traditionally, the responsibility of building out an IT infrastructure has fallen on the shoulders of IT. The IT organization is responsible for not only identifying the organization's business requirements from an IT standpoint, but also determining how to effectively meet those requirements. And one of the most fundamental questions IT must answer is: What kind of storage, servers, and switches does the business need? This crucial question instantly leads to additional questions, which generally relate to features, performance, compatibility, and most importantly, cost.

For large organizations, purchasing responsibility often falls to multiple groups of people, requiring some level of coordination. This adds complexity, resulting in longer lead times, deployment delays, and configuration problems.

DIY is a well-known approach, and most IT organizations are traditionally organized to staff and support this model. But how can this process become better? By layering virtualization onto a 3-tier architecture, you add agility, responsiveness, and better overall utilization to this well-understood and time-proven infrastructure approach. And what's more, individual components can be more easily tuned or repurposed.

### Converged Architecture

A converged infrastructure brings together individual storage, networking, and compute resources into a fully pre-integrated solution that is layered and includes server virtualization. This proven approach has saved businesses many hours designing, configuring, and testing servers, networks, and storage subsystems. The converged infrastructure is delivered as a preconfigured, prebuilt, and pretested solution, which is typically ready to be turned on as soon as it arrives. The converged infrastructure is a hardware-defined solution that can speed provisioning and reduce support and maintenance—although scalability in fixed building blocks can limit flexibility.

### Hyperconverged Architecture

A hyperconverged infrastructure fully collapses the IT stack, delivering all of the infrastructure components in a single server using software-defined capabilities layered with server virtualization. This is possible due to the increased compute power of the latest enterprise-class servers, improved performance and cost of flash storage, and larger capacity spinning drives with smaller footprints. By layering specially engineered software on top of these powerful servers, small-to-medium-sized businesses and distributed enterprises can obtain virtualized infrastructures at a fraction of the cost of DIY infrastructure and converged architectures.

The server turns into an infrastructure building block, complete with virtualized machines, self-contained scale-out storage, and an integrated management interface. Storage cannot be redeployed independently of compute—and individual components can't be repurposed.

Networking is sometimes included in a hyperconverged infrastructure, so it's important to understand the implications of collapsing the stack and routing intra-node communication across the communications network.

In larger organizations, IT is not typically organized to support the hyperconverged deployment model, since it cuts across organizations. This means the approach you ultimately choose should align with your internal IT organizational goals, incorporate legacy infrastructure (from people, processes, and technology standpoints), *and make the shift as seamless as possible*.

In fact, according to ESG research, 88% of IT organizations surveyed are currently using or plan to use converged infrastructure technology solutions and 85% are currently using or plan to use hyperconverged solutions, with cloud infrastructure indicated as a top use case.<sup>3</sup>

That said, when ESG asked IT organizations what they expected their primary infrastructure approach to be five years from now, the majority (57%) answered they would use a DIY 3-tier approach<sup>4</sup>—so it appears likely that on-premises cloud infrastructures will be built using all three types of approaches (i.e., DIY, hyperconverged, and converged).

Lenovo provides infrastructure solutions across the board, and is well positioned to assist IT no matter which approach is used. Lenovo's XClarity systems management software can help IT organizations realize the full benefits of deploying hyperconverged and converged systems by providing a simple way to define, provision, and manage physical infrastructure with software. It integrates into existing IT management applications, such as VMware vRealize and Microsoft System Center, through REST APIs and software plug-ins. And since XClarity is designed for managing the full Lenovo portfolio—including rack, blade, and tower servers; converged and hyperconverged systems; networking; and storage—you don't need separate management stacks for converged, hyperconverged, and DIY environments. Since most organizations run, or will run in the near future, a mix of infrastructure types, systems management software like XClarity will play a prominent role in helping IT organizations run efficiently by minimizing management time and training.

Keep in mind that public cloud can still have a place in modern IT infrastructure. IT professionals often cite public cloud for "burst" workloads as a use case. In reality, most business applications don't need this type of workload scaling because they have predictable scale and performance. Although the same concerns about security and performance still apply, public cloud can be an attractive solution for temporary workloads, such as test and development.

In the real world, many organizations use a mix of public and on-premises cloud. Once you've decided on an infrastructure approach, a number of considerations figure into determining how an on-premises cloud compares with public cloud on a total cost of ownership (TCO) basis. Previously, we stated that deploying an on-premises cloud could be less expensive on an overall TCO basis than moving to the public cloud. Let's take a look at why.

## Buying On-premises Infrastructure versus Renting in the Cloud

At first glance, renting is an attractive option. Benefits include:

- Fast start up.
- Low-cost entry point.
- Pay-for-use pricing with the ability to scale up and down as needed.

For a more comprehensive understanding of these potential benefits, organizations should dig deeper into the cost-to-own discussion.

---

<sup>3</sup> Source: ESG Research Report, [The Cloud Computing Spectrum, Private to Hybrid](#), March 2016.

<sup>4</sup> *ibid.*

## Fundamental Questions for a True Cost-to-own Discussion

- How long before equipment must be replaced? This is generally a set time period (typically covering the three- to five-year lifecycle of IT equipment).
- What will be the capital outlay versus monthly fees?
- How will people and processes be impacted?
- How will space be utilized? Determine the on-premises footprint, including floor space, power, and cooling.
- What will be my capacity requirements? Think virtual machine sizing, CPU, memory, networking, and storage.

## On-premises Solutions

As discussed, IT can take three core approaches to building the on-premises infrastructure. These include:

- **DIY 3-tier.** Server, storage, and network are separately designed, sourced, integrated, and tested.
- **Converged.** One provider makes for easier management, though customization is somewhat reduced.
- **Hyperconverged.** One provider offers integrated appliances, simple systems management (with less customization available), and ease of service and support.

Regardless of the infrastructure option you choose, you will need to take the following costs into consideration:

- **Virtualization software.** Licensing and management from one or more vendors.
- **Installation.** The start-up, sizing, configuring, and cabling of the solution.
- **Maintenance.** Code patches, hardware upgrades, and break/fix.
- **Footprint.** Floor space, power, and cooling.
- **Personnel.** Systems management.
- **Hardware, service, and support.**

## Public Cloud-based Solutions

At first glance, a public cloud-based solution may sound like the ideal option for your organization. You have virtually unlimited processing power allowing you to scale up or down at any time, pay only for what you use, and can potentially reduce internal resources. But with a decision as important as this, you must consider your organization's business needs, and IT and budgetary requirements. You also need to perform due diligence, and take the following factors into account:

- **CPU/core/memory fees.** How much processing power will you need, and what will it cost per month and over the long term?
- **Storage fees.** What are your anticipated and projected storage capacity and performance needs, and how will storage fees add to your monthly bill?

- **Bandwidth charges.** Ingress charges are typically built into the core subscription price. While it's easy to get data in, egress fees may apply if significant amounts of data are pulled out.
- **Personnel costs.** The manpower costs associated with managing the cloud environment, applications, IT integrations, and vendor relationships need to be considered.
- **Utilization.** Reserve enough monthly CPU, memory, and storage to accommodate spikes (since billing is monthly and not real time).
- **Network performance.** You pay for the level of processing and performance you need—not just for individual components, but for network performance as well. Keep in mind that network performance can be a huge challenge because you can't necessarily control the amount of hops between every machine you're using.
- **Risk and compliance.** The complexity, risks, and costs associated with regulatory compliance and reducing the risk of a cyber-attack must be top of mind.

While it's clear that a public cloud-based solution can simplify IT administration once it is up and running, some thought needs to go into these considerations based on the needs of your business. That said, the bottom line consideration—and the primary focus of this report—is the true cost of public versus private cloud over time.

## Comparing the Cost of Public Cloud and Private Cloud Powered by Intel and Lenovo

In this section, we'll be exploring the results of some economic modelling with a goal comparing the cost of renting virtual machine infrastructure from an industry-leading public cloud vendor versus purchasing and managing private cloud infrastructure powered by Lenovo with Intel inside. The Enterprise Strategy Group has built a number of models that compare the costs of industry-leading public cloud vendors to the cost of converged and hyperconverged private cloud infrastructures. As you'll see in this section, when you level the playing field and consider all of the costs over a typical period of ownership (i.e., three years), the total cost of ownership for a private cloud is often more affordable over time.

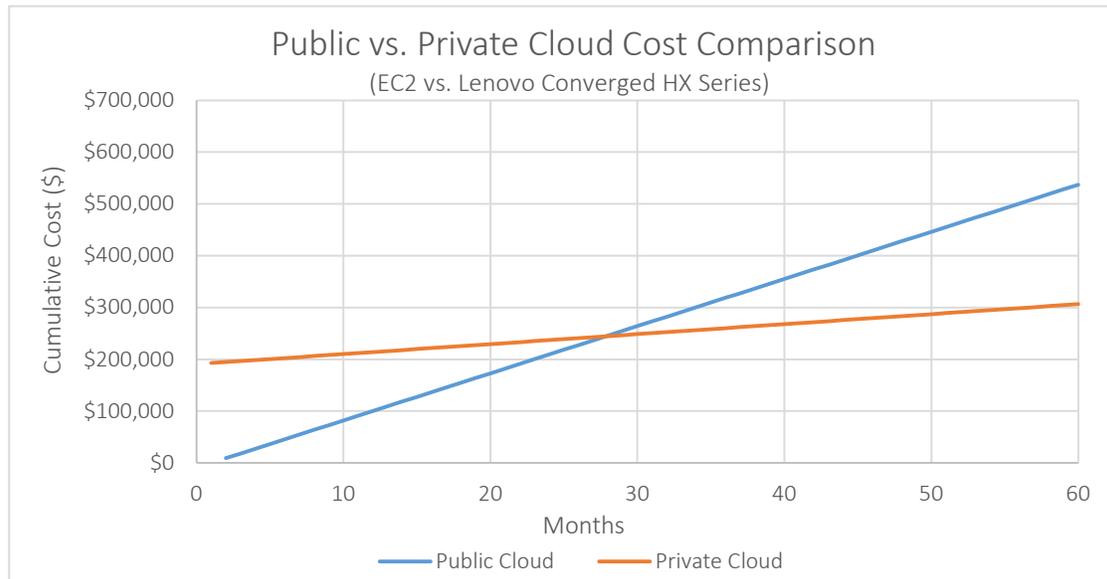
For this report, the publically available pricing of an industry-leading public cloud vendor was compared with the price of a private cloud that was built with a four-node Lenovo Converged HX-Series cluster. A virtual machine infrastructure was modelled with a goal of assessing the costs of an infrastructure that can support 160 medium-sized virtual machines.<sup>5</sup>

As you can see in Figure 2, the cumulative rental cost of the public cloud option is more affordable at first, and is under \$100,000 for the first 12 months. Those costs add up over time. The cost of the private cloud option, which includes the upfront capital equipment cost of the private cloud hardware (server, storage, and networking) and infrastructure software, is more expensive at first. The variable costs associated with maintenance, power, cooling, space, and management manpower add to the cost of the private cloud option over time. The important factor to notice here is the crossover point where the cost of the private cloud is less than the cost of renting a public cloud—in this case, it occurs after about two years (28 months). After the crossover point, the savings start to add up. Over longer periods of ownership, the savings can be substantial—in this example, the Intel Lenovo private cloud is nearly a quarter million dollars (\$230,693) cheaper than the public cloud over a five-year period of ownership.

---

<sup>5</sup> Lenovo Converged HX3710, four nodes in 2U with 24 cores per processor; 160 VMs, each with 4 GB of RAM and 500 GB of storage capacity (26 TB total) versus [publically available](#) AWS EC2 pricing as of January 13, 2017: Linux on t2.medium VMs (2 vCPU, 4 GB of RAM), on-demand pricing, EBS HDD storage.

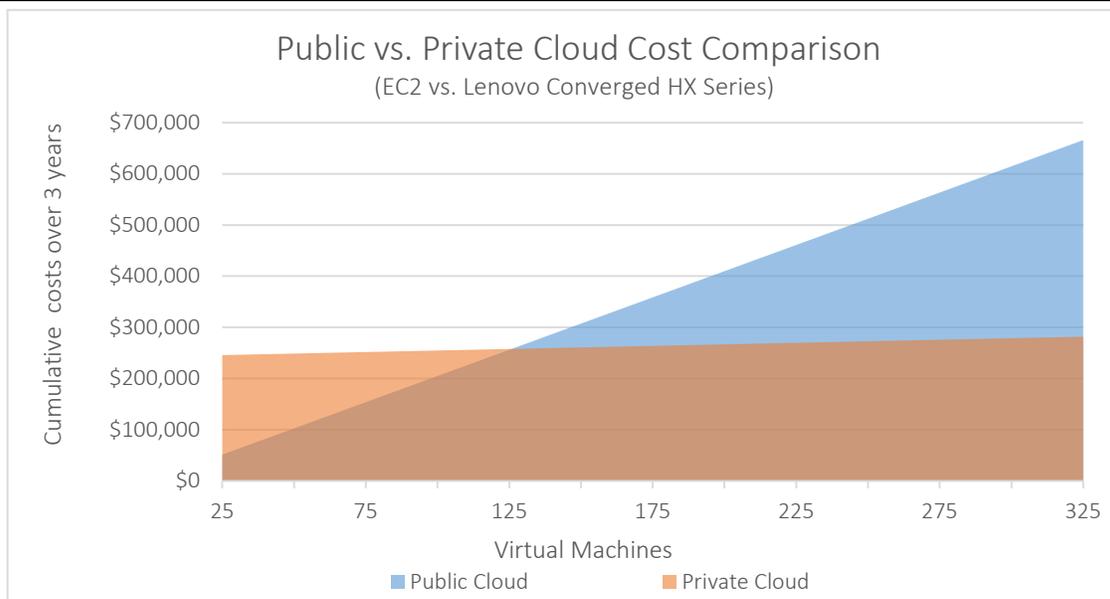
**Figure 2. Public versus Private Cloud Cost Comparison Over Time**



Source: Enterprise Strategy Group, 2017

We also examined the three-year cost of ownership of the same Lenovo private cloud infrastructure for a growing number of virtual machines. Once again, as you can see in Figure 3, the cost of owning the private cloud is higher at first, but crosses over the cost of the public cloud over time and as the environment scales. In this example, the cost of a private cloud is less than the public cloud for environments with 125 or more virtual machines. As the environment scales, the savings become more pronounced. At 325 VMs, the private cloud is less than half the cost of the public cloud (58% less).

**Figure 3. Public versus Private Cloud Cost Comparison – Virtual Machines**



Source: Enterprise Strategy Group, 2017

While the results shown in Figure 3 provide a clear illustration of private versus public cloud costs for different numbers of virtual machines, the fixed four node Lenovo Converged HX-Series cluster that was used for this analysis is obviously over-

provisioned for smaller numbers of virtual machines (e.g., 25-50). In the real world, an organization could buy fewer, less powerful nodes to start and would add nodes over time to accommodate growth. The general concept shown in the graph still applies, but it would start with a lower price point that grows over time in a step function as nodes are added to meet the needs of the business. It should be noted that one of the great economic benefits of a Lenovo Converged HX-Series cluster is the simplicity of adding right-sized nodes to meet the growing needs of the business.

While the cloud economics shown in Figure 2 and Figure 3 are specific to one private cloud deployment option and one leading public cloud provider, a similar methodology can be used to evaluate the economic crossover point for a wide variety of Lenovo private cloud hardware options including servers, storage, and converged systems. That said, the crossover point where private cloud savings are achieved, and the magnitude of the potential savings, depends on a number of factors that are specific to your organization. In other words, your mileage will vary depending on a number of factors including:

- The performance and capacity requirements of your virtualized applications.
- The number of virtual machines over time.
- The cost of manpower to maintain the infrastructure.
- The cost of virtualization software and the manpower to maintain that virtualization software (note that Lenovo hyperconverged solutions have optional, built-in virtualization software at no additional charge, and are known for industry-leading ease of management).
- The costs associated with data protection and disaster recovery.
- The potential costs associated with public cloud WAN egress.

The economics of cloud are similar to leasing or buying a car. After the initial costs of purchasing a car (or paying off a car loan), the cost of owning a car becomes less than the cost of leasing a car over time. That crossover point depends on a number of factors, including how long you plan on driving the car, miles per gallon, the number of miles you typically drive per year, the cost of money, etc.

## The Bigger Truth

IT needs to build for business agility to keep up with the constant state of change. Adopting cloud technologies, whether from a cloud service provider or by building a private cloud, can help. The public cloud offers fast startup and low-cost entry points. Utilizing public cloud is a tempting “quick fix” for building an agile IT environment, reigning in IT costs, or starting up a new service. However, it may not be a cost-effective or long-term solution for many applications.

The public cloud has an unquestionable place in the IT world. It can be a great solution for transient workloads, test and development, and temporary dynamic workloads (such as ad campaigns). It is also affordable as a backup or disaster recovery target, since it provides offsite capacity without IT needing to build out remote infrastructure—a use case that drives significant cost savings on its own. But for *known* workloads (such as those found in most business application environments), the cloud may not live up to its “cost-effective” reputation. In fact, many IT organizations have been surprised by how their monthly bills grow due to egress fees, the cost of scale (VMs, CPUs, and memory), and price premiums for better performance (memory or bandwidth). For this reason, many organizations are taking a hybrid approach, using the cloud for some workloads while keeping known workloads or those dealing with sensitive data on-premises.

IT’s objective isn’t to choose a winner between cloud and on-premises—it’s about finding what works for your business from agility, risk, and cost standpoints. The answer to these concerns is using both private and public, often in a hybrid fashion. Advances in technology, such as converged and hyperconverged solutions, take much of the complexity, risk, and management overhead out of the equation by collapsing the stack into fewer components managed in a software-defined fashion, such as those available in the Lenovo solutions. This technology makes building a private cloud a viable, cost-effective (especially over time and at scale) alternative to add agility and reap benefits similar to those that come from using the public cloud. The “what works for the business” discussion is now about understanding the cost, risk, and agility tradeoffs for each workload, and placing the workload where it makes sense—either on-premises or in the cloud.

Cloud technology is advancing quickly. While server virtualization is a mature market, technologies like converged or hyperconverged systems that take server virtualization to the next level to support building private and hybrid clouds are still relatively new. And many vendors are working in the “software-defined” space to drive more agility, simplicity, and value. It’s important to note that Lenovo does not have an enterprise software unit—instead, it partners to provide leading software-defined solutions to its customers. This allows Lenovo to work with vendors across the IT spectrum to find the best solutions to meet customer needs and bring the most advanced software technology to market quickly. This practice will continue to drive value up, and costs down. The economics can only improve from here.

As the models show, a private cloud can be a more cost-effective option than the public cloud in just over two years. The cost savings accelerate with scale, but can quickly reach the six figure (or more) range. Recognizing that IT environments differ, it is up to IT organizations to do the research and explore their anticipated needs.

***Do the homework.*** It’s imperative for IT decision makers to investigate different infrastructure approaches for building an on-premises cloud, and compare these approaches with the cost of public cloud *as you scale and over time*. Only then can the breakeven points and business benefits be fully understood so workload placement can be determined.

ESG encourages you to connect with your Lenovo representative for assistance in performing a cost analysis for your particular needs.

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.



**Enterprise Strategy Group** is an IT analyst, research, validation, and strategy firm that provides actionable insight and intelligence to the global IT community.

© 2017 by The Enterprise Strategy Group, Inc. All Rights Reserved.

