White Paper

Seeing Past the Hype: Understanding Converged and Hyperconverged Platforms

By Terri McClure, ESG Senior Analyst

November2016

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

Contents

Background: The Emergence of Convergence

There is a fundamental change underway in how we run IT—a shift in infrastructure deployments to allow for faster response time, greater scale, manageability, simplicity, and an overall better service and support experience. It's not done in a vacuum—it's for a good reason: to accelerate business processes. Many businesses recognize that of all the organizations within an enterprise, IT is the one that can invariably make the most difference when leveraged optimally and strategically. Speeding access to applications and information can help businesses quickly take advantage of new opportunities, accelerate time to market for products and services, and increase workforce productivity.

At a high level, businesses are looking for their IT organization to run in a more cloud-like manner, operating efficiently and with the agility to spin up applications in minutes rather than days and weeks. This requires a shift in how IT architects the underlying infrastructure to provide a more services-like experience for stakeholders—whether they are employees, customers, or prospects. A major instantiation of this shift at an infrastructure level is in the shift to adoption of converged and hyperconverged platforms.

Why convergence? When ESG asked those organizations that have adopted or plan to adopt converged or hyperconverged platforms their reasons for considering these platforms, the actual or expected benefits make a clear case that this infrastructure approach helps IT progress toward cloud transformation. The most cited primary reasons for adopting converged or hyperconverged solutions include improved service and support, increased agility in virtual machine provisioning, and better speed of deployment (see Figure 1),[1] all keys to helping stay agile and always-on. But organizations are also looking at these platforms to achieve scale with predictable costs and to ease the infrastructure management burden, among other reasons.

Figure 1. Reasons for Considering or Deploying Converged and Hyperconverged Platforms

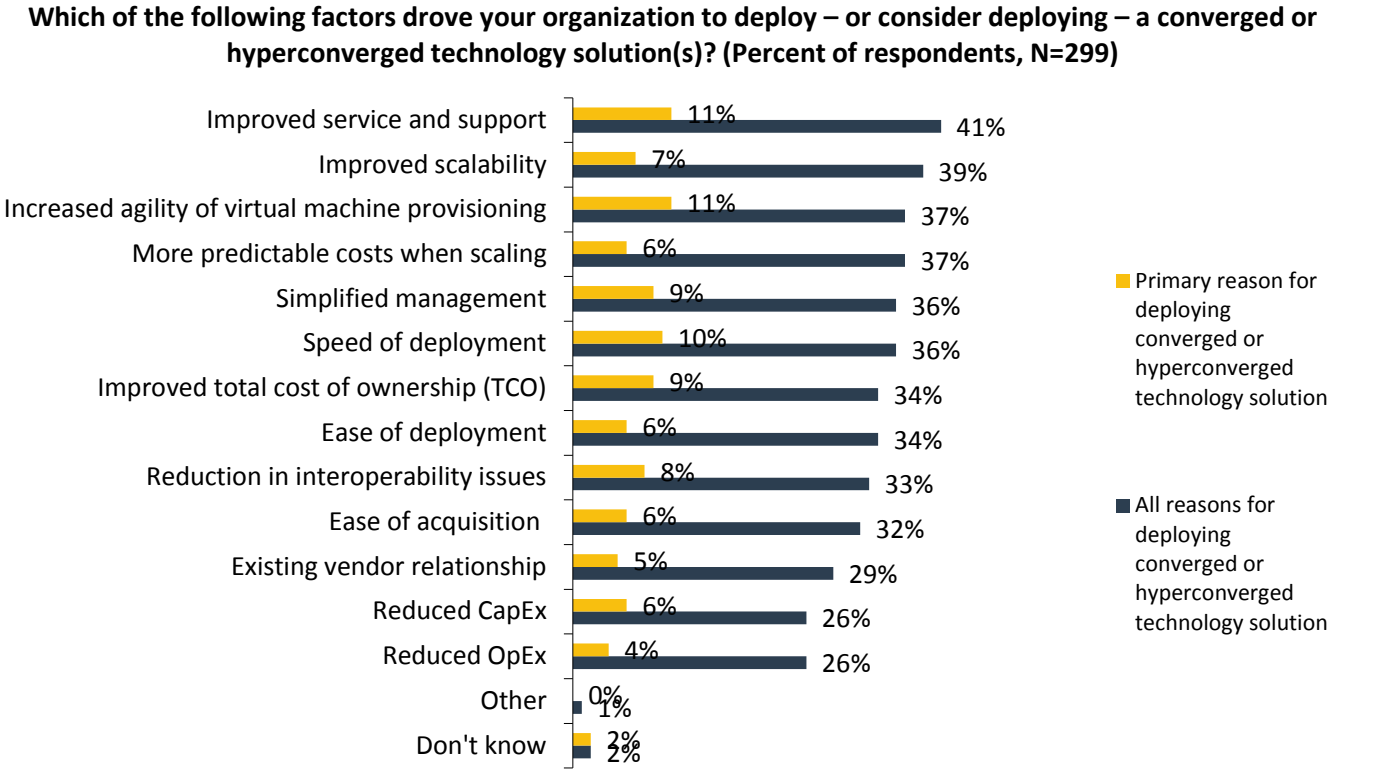

Source: Enterprise Strategy Group, 2016

[1] Source: ESG Research Report, *The Cloud Computing Spectrum, from Private to Hybrid*, March 2016. All ESG research references and charts in this white paper have been taken from this research report.

As you study the actual and perceived benefits IT organizations are citing for adoption, it should come as no surprise that a clear majority of IT organizations surveyed use or plan to use converged and/or hyperconverged systems (see Figure 2).

Figure 2. IT Usage of and Plans for Using Converged and Hyperconverged Platforms

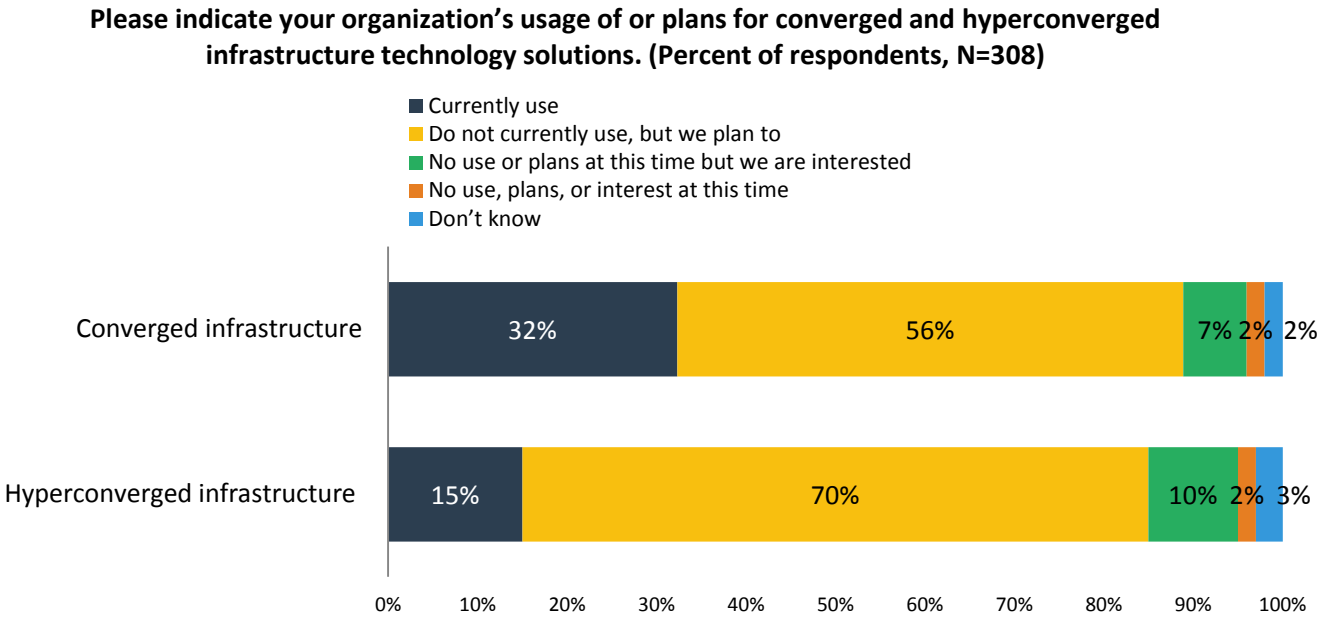

Source: Enterprise Strategy Group, 2016

Given all the buzz about convergence and hyperconvergence in the enterprise, it's good to understand just what these systems are, how they differ from more traditional approaches to IT, and how they help with the much longed-for cloud transformation in the data center—as well as the pitfalls to each approach.

Understanding Converged and Hyperconverged Platforms

It is important to understand the starting point. Traditionally, the responsibility of building out an IT infrastructure has fallen on the shoulders of IT. They are responsible for not only identifying what their business requirements are from an IT standpoint, but also figuring out how to meet those requirements. And the big question they must answer is, "What kind of storage, servers, and switches do I need?" That question instantly leads to many others relating to features, performance, compatibility, and, most importantly, cost. For large organizations, the purchasing responsibility often falls on multiple groups of people, requiring some level of coordination. This adds complexity, resulting in longer lead times, deployment delays, and configuration problems.

But IT has managed because essentially it had no choice but to proceed that way. And it did well, even though each data center was essentially buying car parts and building their own car, sometimes with a lathe and mallet to make it all work! That cost money—whether for internal staff (fixed cost) or consultants (variable but even more expensive)—and time either way. Things had to change, and that's where the first shift in consumption models started, essentially with the move to reference architectures (RAs), where at least IT knew it had the correct car kit and some instructions!

RAs offer blueprints of a clearly defined set of compatible software and hardware for building out IT infrastructures. By following the blueprints and best practices outlined by their desired vendors, IT administrators can confidently build out a solution that meets their requirements, whether it's for a specific, mission-critical application or for a mix of applications used by everyone in the business. RAs have paved the way for the rapid adoption of virtualization for those organizations that have not yet made the leap. Although RAs help address *some* of the complexities of the build-your-own-infrastructure (BYOI) approach, specifically around interoperability, issues still exist because of multiple support points and a lack of

comprehensive management synergy between components. Converged and hyperconverged systems developed as architectural approaches to reduce complexity even further.

First There Were Converged Systems

Converged infrastructures build upon the RA idea, but take it a step further by assembling and testing the infrastructure prior to shipping. Converged solutions take individual storage, networking, and compute resources and bring them all together into a fully pre-integrated solution. This approach has proven to save businesses countless hours designing, configuring, and testing servers, networks, and storage subsystems. The converged infrastructure gets delivered as a pre-built, pre-tested, and pre-configured solution that is usually ready to be turned on as soon as it arrives. And because the components are integrated at a hardware level, repurposing an individual component like a server or switch is easy, and most of the resiliency comes from the hardware (hardware redundancy).

Storage management is a key point of differentiation between converged and hyperconverged architectures. In a converged system, storage is managed much like DAS—configuring LUNs, volumes, and RAID groups that are integrated with the server and network. Because this is a best of breed system, typical storage features are supported—snapshot, clones, and remote replication. Any features found in the enterprise-class array offered with the solution remain available in the integrated system, providing investment leverage and protection of existing business processes.

A wide range of converged configurations is available, all designed with workload support requirements in mind to determine proper CPU/memory/storage ratios to meet application performance requirements.

Benefits of converged systems:

- It is a mature technology, well understood by IT professionals.
- IT can benefit from known and predictable performance.
- Hardware can be shared or repurposed.
- A converged system is supported by a single vendor.

Then Came Hyperconverged

Hyperconverged infrastructures are software-defined appliances that combine server, storage, virtualization, and (sometimes) networking resources into a software-defined building block. This is made possible due to the increased horsepower of the latest enterprise-class servers, improved performance 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 can implement virtualized infrastructures at a fraction of the cost of a 3-tier architecture. The server turns into an infrastructure building block, complete with virtualized machines, self-contained scale-out storage, and an integrated management interface.

Hyperconverged solutions scale by adding new appliances that combine into a cluster, essentially creating a shared resource pool. And while convergence relies on tightly integrated hardware, hyperconvergence relies on software to knit the components together, and resiliency comes from the software. Hyperconverged solutions are fully virtualized across servers, storage, and networking. As such, resources are managed from the top down—storage management is typically based on VM-defined virtual volumes, rather than LUNs and RAID groups.

Benefits of hyperconverged solutions:

- They are typically appliances that are fast and easy to deploy.
- The entry point is low impact from both capacity and cost standpoints.

With hyperconverged solutions, the entry point is attractive, but with some systems it is difficult to get linear scalability with this type of building-block approach. Essentially, as nodes are added, there is a set amount of both CPU and storage capacity added—likely more than users need of one or the other. Several vendors have introduced either storage- or CPU-intensive nodes to help provide better balanced scalability.

Because the hyperconverged market is not as mature as the converged market (albeit it is rapidly maturing), there are some other considerations associated with many vendor's solutions. Host side integration is one area that still needs development (e.g., VMware vvol support, VSS plug-ins, and Microsoft SCOM plug-ins). Some of that exists in the hyperconverged landscape, but for many vendors these integrations are still underway. Best practices and integration at a higher application level is just beginning as well (e.g., Microsoft SQL Server).

What Does All This Really Mean?

There is a lot of market hype around converged and hyperconverged infrastructures—so it is important to understand how to decide what is optimal for a given workload or situation. By the way, by "hype" we mean that adoption—while fast and growing, make no mistake—still trails what you'd expect from reading the press and blogosphere. Indeed, more traditional 3-tier systems still outweigh the converged and hyperconverged world (see Figure 3). However, there is clear value and convergence generally is emerging.

Unfortunately, determining the optimal approach to any given workload is not simple. Like almost everything else in IT, there are many considerations that go into the equation. And the best fit comes down to many things, *not the least of which are organizational, cultural, and practical.* Shifting without considering many factors can be risky. A tremendous amount of time and money has been invested in the people, processes, organizational structure, and technology that support the business today. The paramount consideration for IT should be "first thou shall do no harm"—keep applications up and running, and data protected and secure. This is why the majority of workloads will be supported with traditional 3-tier architecture in the near term (see Figure 3). It's what IT organizations know and are organized to support, and there are tried and true practices around managing and protecting data in this type of configuration.

Figure 3. IT Infrastructure Strategy, Five Years Out

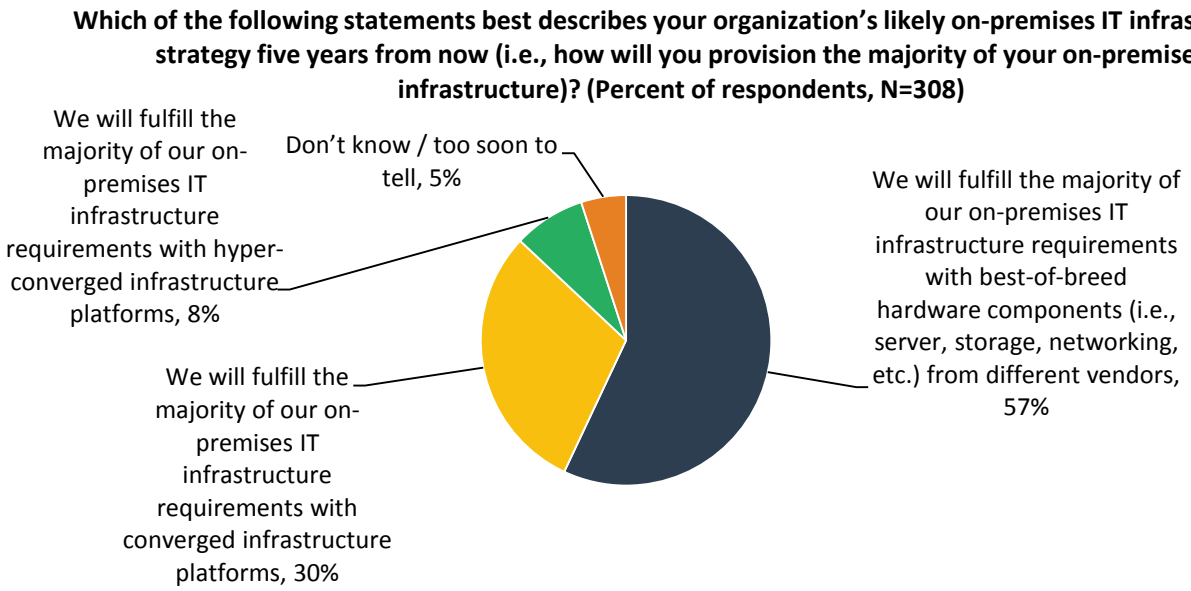

Source: Enterprise Strategy Group, 2016

The View from Current and Planned Converged and Hyperconverged Users

That said, and as previously illustrated, the converged and hyperconverged movements are happening, and nearly 40% of the IT pros we asked about future architectural plans are moving toward these platforms in a major way. So what do today's IT organizations that are using or evaluating converged or hyperconverged solutions think? How do they determine the best architectural approach?

One of the fascinating things in IT is that what users believe/perceive clearly shapes their decisions. ESG research shows why users say they have been making their CI and HCI solution decisions so far based on beliefs and perceptions (see Figures 4 and 5).

Figure 4. Reasons for Using Converged over Hyperconverged Platforms

Source: Enterprise Strategy Group, 2016

Early adopters and current users provided several reasons for preferring converged over hyperconverged solutions. However, performance, reliability, and scale top the list. For these users, a hardware-based approach provides an ability to achieve predictable performance and isolate workloads onto separate hardware resources—essentially allowing for more control over how resources are allocated. This explains why many organizations believe that CI solutions provide a better approach to supporting tier-1 workloads compared with HCI.

Figure 5. Reasons for Using Hyperconverged over Converged Platforms

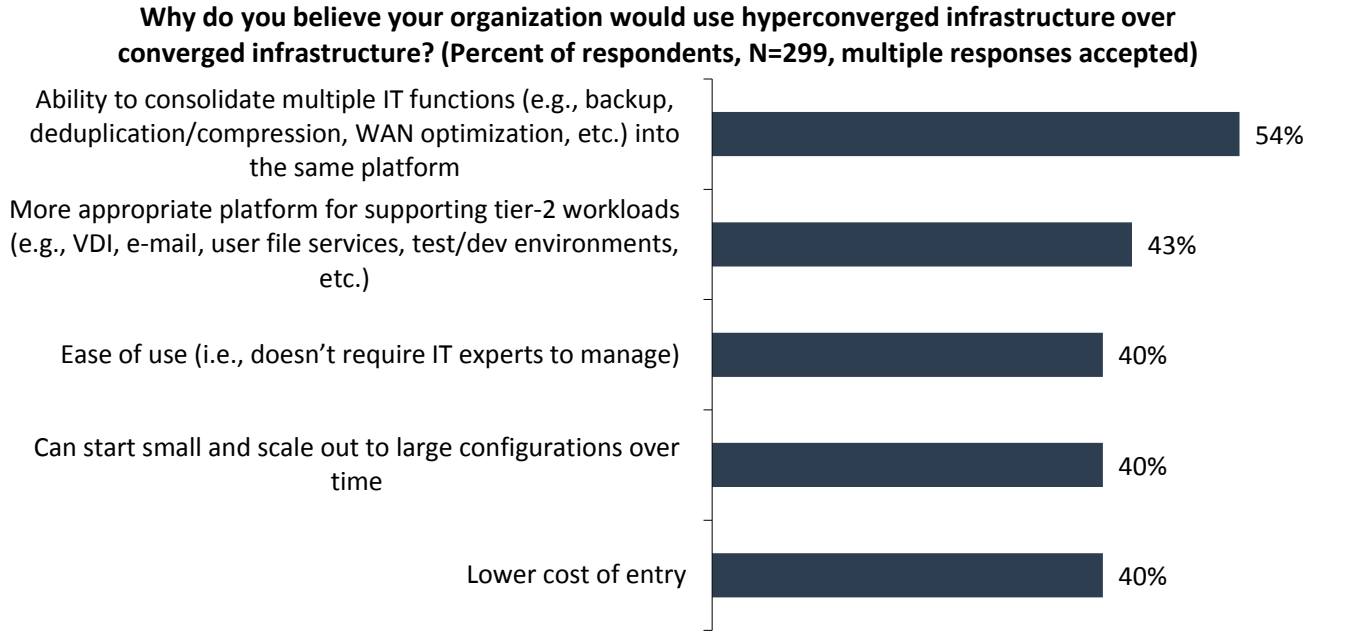

Source: Enterprise Strategy Group, 2016

And of course, users also provide several reasons for preferring HCI over CI solutions, again whether real or perceived. Given the software-defined, appliance-based, building-block approach these solutions provide, users feel that they are ideal for resource consolidation, tier-2 workload support, and ease of use. Consolidation comes from both workload and hardware consolidation, eliminating the 3-tier architecture. And these appliances come with an attractive price point compared with 3-tier and converged approaches, delivering economy through the software-based architectural management and redundancy.

So at the end of the day there are benefits to both approaches—they've each found a niche, for now. Ultimately, it often comes down to perceptions about scale and price.

Networking and License Implications: The Potential Hidden Costs of Hyperconvergence

Due to the building-block, clustered appliance approach used in hyperconvergence designs, it is important to understand some broader implications of deploying these systems, particularly when it comes to networking and software licensing. Certainly, these issues won't apply to everyone or every use case, but IT organizations need to be aware of them as they evaluate solutions.

In hyperconverged systems, intra-cluster communication and storage traffic are routed through the communications network that serves the broader organizational requirements. Using some hyperconverged solutions for tier-1 applications that require a large amount of intra-cluster communication or for applications that require large data sets may be impractical. Of course, hyperconverged solutions can be used for these types of applications, but they may require investing in a redundant network to handle traffic, negating some portion of the savings from deploying them.

The cluster-based approach also means that licensing for some tier-1 applications may need to be renegotiated. With hyperconverged solutions, you often need to add processor nodes in order to add storage capacity, which means more CPU cores are added. Traditionally, companies like Oracle and SAP charge per core, so as hyperconverged systems scale and appliances are added, the core count adds up. This is especially important if only some of the nodes are running Oracle or SAP software. Some of the hyperconverged vendors have ways to either contain these applications to a subset of

nodes/resources within a cluster, or even partition, but this needs to be an upfront consideration so as to avoid surprises when licenses must be trued up.

The Bigger Truth

Many factors come into these decisions—budgets, "religious" beliefs about IT, satisfaction with current vendors, and support…and *most important of all,* what it is that you are actually trying to achieve: Users should start with the desired outcome and ask their chosen/considered vendors for one or more possible routes to achieve that outcome. Decisions should be made based on what best meets the prescribed criteria and promises the optimal required result, irrespective of its semantic nicety. In IT as with clothing, fashion is fleeting—comfort, style, and fitness for purpose don't age.

Converged is a more mature/configurable option that helps organizations customize configurations to better support tier-1 apps via predictable performance—today. Converged allows IT organizations to repurpose hardware or easily distribute resources from the whole stack to a huge organization. It's harder to do that today with hyperconverged appliances.

Things will change over time, and the market bears watching. All the major vendors offer converged solutions, such as the FlexPod solution offered by NetApp—it is a more mature approach. And while FlexPod may not suit every situation, certainly when converged platforms are the optimal solution, FlexPod is a well-known and reliable choice. We expect all the big vendors to offer both options. Like SAN, NAS, and DAS, they will simply become IT building bricks and not mutually exclusive offerings/choices. Hyperconverged is largely, though not completely, the domain of startups today (several of the major vendors with hyperconverged solutions OEM them from startups). But don't expect the major vendors to maintain the status quo. Portfolios will expand, hyperconverged solutions will continue to mature to address a wider variety of workloads, and it would (of course!) make sense for vendors like NetApp and others to get on the hyperconverged bandwagon.

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

 www.esg-global.com

 contact@esg-global.com

 P. 508.482.0188