

Hyperconverged Infrastructure: Improve business value while decreasing TCO

White Paper

Author: Benjamin Woo



Neuraspective™

Key Findings

Today's modern enterprises need to reexamine the role technology plays within their organizations. Many still look at technology as a cost, rather than an investment. Traditional metrics such as ROI have become meaningless when deploying highly virtualized and multi-tenanted information infrastructure.

Instead, enterprises need to consider new metrics. In particular, total cost of ownership (TCO), net increase to enterprise value and competitive advantage are more appropriate metrics for today.

At the same time, information infrastructure, like software, needs to be able to be acquired and provisioned in a predictable, repeatable and scalable way.

Hyperconverged infrastructures (HCI) enables enterprises to lower their TCO, improve agility (leading to the ability to quickly generate more growth) and the ability to focus on data and information, as opposed to infrastructure.

In this White Paper, Neuralytix looks at what is HCI, and in particular a leading provider, Gridstore's approach to HCI.

Contents

Overview	5
What is Hyper-Converged Infrastructure (HCI)?	6
Why take a CI or HCI approach at all?	8
Who Is Gridstore?.....	9
Why Hyper-V?	9
Why Gridstore?.....	10
HCI Considerations	12
Not all integrations are alike	13
The Gridstore Difference	13
Per VM Performance, Capacity and QoS Control	14
Density without sacrificing reliability	15
Independent Scalability.....	17
Conclusion	18
About the Author	19
About Neuralytix™.....	19

Analysis

Overview

Neuralytix forecasts that starting in 2015, enterprises, like service providers have been doing, will increasingly look to converged infrastructure (CI) and hyperconverged infrastructure (HCI) to satisfy their datacenter needs.

The reason is simple. HCI has the ability to drive down the total cost of ownership (TCO).

In today's information-driven world, the key to generating business value and competitive advantage is maximizing IT agility while minimizing TCO.

The old metrics such as return on investment (ROI) or driving down capital expenditures are no longer applicable. Relying on these outdated metrics can actually hurt the modern enterprise more than help it.

The capital cost of datacenter equipment can represent as low as 20% of the TCO. So optimizing on capital expenditure has a relatively low impact for an enterprises. In fact, for some do-it-yourself (DIY) enterprises, the cost of maintaining, supporting and servicing low capital cost infrastructures and architectures can dramatically increase the TCO.

With the virtualization of almost all the infrastructure within a datacenter, ROI metrics have absolutely no place in the technology decision. It is impossible to measure the ROI on any piece of



equipment, given the seemingly limitless and “on-demand” ways in which any given piece of equipment can be deployed.

What becomes important in this equation is performance and data capacity. Optimizing performance and capacity allows enterprises to create information agility. This in turn allows the enterprise to create business value and competitive advantage.

Specifically, HCI has the benefits of:

- Simple design
- Predictable performance
- Minimal power and cooling
- Lowered management costs
- Reduced IT admins required
- Scalable
- Repeatable deployment
- Best total cost of ownership (TCO)

Moreover, HCI allows IT departments to focus on *Information* rather than *Infrastructure*, and as such, enable IT departments to help their respective organizations improve their company’s business value and competitive advantage.

[What is Hyper-Converged Infrastructure \(HCI\)?](#)

HCI is where the five key components of a datacenter are pre-integrated and packaged in a single form factor. This single form factor is then sold, serviced maintained as a single modular unit.

The five components of the datacenter are:

- Computing,



- Virtualization/Hypervisor,
- Networking
- Data storage, and
- Datacenter management software

HCI takes the concept of converged infrastructure (CI) and adapts it for minimal footprint, maximum standardization and modularization, and optimizes it for the lowest TCO and the simplest management.

By creating a standardized datacenter module, the following areas can be minimized:

- Management
- Power and cooling
- Number of administrators needed

This is key to minimizing TCO. By standardizing all of these components into a single modular unit, customers benefit from:

- Less rack space
- Less power and cooling
- Single management across all modules
- Reduce number of IT admins
- Significantly reduced IT design costs

The last bullet point above is key. HCI is ***not*** a do-it-yourself (DIY) approach. It is a very considered logical approach to a standardized datacenter.

Why take a CI or HCI approach at all?

CI and HCI help solve a key problem: improve business value. In today's data driven business world, IT departments must focus on driving value from data, and spend less time on the datacenter and its infrastructure.

After all, IT stands for *Information* Technology, not Infrastructure Technology!

Optimizing infrastructure, particularly in today's highly virtualized, highly complex datacenter can be grueling. In particular, optimizing performance across multiple workloads of varying I/O characteristics residing on any given server causes Quality of Service (QoS) issues that are often very difficult to find and solve.

Then there is the management of new technologies such as flash. While flash can help to improve performance, it can only help if it is properly deployed. Overprovisioning flash is at least as expensive as overprovisioning any other storage medium, if not more, and still may not help to solve specific performance issues!

Who Is Gridstore?

Gridstore is a leading provider of hyper-converged infrastructure purpose-built for Microsoft workloads. Gridstore's HCI platform is built on patented software that is massively scalable, simple to deploy and easy to manage.

As noted earlier, HCI combine all the infrastructure components of a datacenter into small form factor. A common misnomer is that cost is the primary driver for most IT organizations. While cost is a *major* driver, most organizations are likely to pay a (reasonable) premium for a simpler-to-manage system.

The rationale is simple. The costs of additional IT admins or a more skilled admins, including benefits, etc. outweighs the marginal premium of an IT solution that is already simpler!

One key differentiator to Gridstore's success is its sole focus on the Microsoft Hyper-V platform.

Why Hyper-V?

First, there is cost. Hyper-V is free with the latest versions of Windows. Although, it can be argued that the VMware ESX hypervisor and KVM are also essentially free, the ancillary costs of plug-ins, integration and management software can quickly add up.

Given that the majority of applications are developed for Windows, this makes Hyper-V a perfect hypervisor platform.

Over the last several years, Hyper-V has come into its own, not because it is free, but also because it is ready for prime time. Many

large enterprises have standardized on Hyper-V and the market adoption of Hyper-V has significantly increased.

Hyper-V is now the underlying hypervisor for many mission critical deployments of virtual desktop, Microsoft SQL Server, Microsoft Exchange, etc.

Why Gridstore?

Gridstore's HCI solution meets the three key criteria for today's information infrastructure:

- Predictable;
- Repeatable; and
- Scalable.

Predictability when it comes to today's information infrastructure is arguably the most critical criterion. Every infrastructure component should be predictable in performance and capacity. This is the only way in which the next generation (or more accurately, the new generation) of highly virtualized datacenter can be effective.

Gridstore' ability to deliver end-to-end predictable performance and QoS on a per VM basis makes it a leading contender for any HCI consideration.

As well as being predictable, HCI should be repeatable. This simply means that each time there is a new requirement for additional resources, that the implementation, administration and support of new resources should be identical. The repeatability aspect of HCI continually drives down the TCO. Gridstore's ability to add



incremental nodes as needed allows for predictable, repeatable scaling over time.

Finally, all HCI should be scalable. Scalability should come in the form of scale-up and scale-out.

Scale-up capability would be the ability to add more powerful, or more capable resources in the same footprint, while scale-out capability would be the ability to seamlessly add the same resources and view the holistic “grid” or cluster as a single logical entity. With Gridstore you can start with as few as three nodes and scale to a total of 250.

Neuralytix believes that the Gridstore solution meets all three (predictable, repeatable and scalable) of these essential criteria.

HCI Considerations

With HCI, there are only two main things to worry about – performance and storage capacity.

Both of these factors can be easily solved through scale-out technologies. Scale out technologies enables the user to buy pre-configured, pre-integrated modules, irrespective of the capacity, and then build an infrastructure out to the desired scale. Consider the concept of the basic Lego 4x2 brick. Irrespective of color, the 4x2 brick is still the same. So long as you have enough of them, you can build pretty much anything. This applies to a scale-out architecture also.

The huge benefit from this approach is that every “brick” or every system in the datacenter is essentially identical. Every IT admin (if there is more than one) will know exactly how to manage every component in the datacenter. Knowledge can be easily shared. This reduces the need for specialized IT administrators, perhaps reducing the number of IT administrators needed, if budget becomes constrained.

As mentioned earlier, Gridstore focuses its solutions on Microsoft Hyper-V and Windows. This allows Gridstore to work closely with Microsoft to create very tight integration.

This integration allows Gridstore to create a number of interesting and important technical differentiations that make its solution unique.

Not all integrations are alike

When it comes to hypervisor integration, not all integrations are alike.

For example, with VMware, its partners are unable to tap into the kernel. This dramatically limits the ways in which optimization can take place. In most cases, “optimizations” is performed through software loaded onto a virtual machine.

In VMware, performance and I/O optimizations are simply not efficient. It requires the use of virtual machines that talk to other virtual machines via virtual tunnels, as opposed to being able to deploy optimizations directly to the hypervisor itself.

This does not allow true optimizations to take place because everything happens outside the ESX kernel. At best, these “optimizations” are best practices, but in no way improves the operations of the underlying hypervisor.

The Gridstore Difference

On the other hand, Microsoft allows Gridstore to work at the kernel level, in a way that VMware and other hypervisors do not. This means that Gridstore can optimize its capabilities from inside the kernel.

This is unique to Gridstore, and gives it the high performance characteristics and high data reduction rates it achieves.

While hypervisors are designed to be generic in nature, Gridstore is able to truly optimize the way Hyper-V interacts with the

applications that are expected to run within the Gridstore Hyper-V environment.

The result includes unique features such as:

- Granular VM by VM optimizations
- Ability to provide true QoS on a per VM basis
- Using standard Window System Center monitoring and management software.
- Leverage Windows embedded capabilities including
 - VM Snapshot, Replication, Failover/failback, and Live Migration
 - Deduplication
 - Storage Live Migration
 - Thin Provisioning and Trim

Per VM Performance, Capacity and QoS Control

The highly virtualized environments in which we operate demands that each physical server runs multiple workloads. Depending on what each workload requires, the performance desired will differ.

Traditional “one size fits all” approaches does nothing to optimize workload requirements.

To overcome this, Gridstore uses a patented Server-side Virtual Controller Technology (SVCT). These controllers, called vControllers, are able to recognize and detect I/O patterns on both the host and guest machines.

The SVCT is then able to automatically self-optimize the performance for each workload. The self-optimization technology

means that it minimizes the administrators' time required to ensure the optimal performance of the entire environment.

For those applications that need a sustained and guaranteed level of performance, Gridstore has a technology called TrueQoS. Unique to Gridstore, this allows the customer to deliver end-to-end storage quality of service (QoS). This sets the priority level for I/Os at the application server or VM level, and enforces it throughout the entire Gridstore infrastructure.

Neuralytix believes that Gridstore's ability to deliver self-optimizing performance, as well as per VM QoS, distinguishes from its competitors.

[Density without sacrificing reliability](#)

One of the many benefits of HCI is the density at which compute and data resources can be delivered using common off-the-shelf (COTS) elements to drive down capital costs.

However, density cannot come at the expense of reliability.

Some approaches to this include the "triple mirror" technique, in which all resources are replicated three times. By doing so, the cluster or "grid" can sustain multiple component failures without affecting the uptime of the entire datacenter.

Triple mirroring is costly – it is at least three times the cost!

Gridstore has cleverly deployed modern data protection technologies such as erasure coding to overcome this cost. In doing so, Gridstore can lower the TCO by as much as 50% over its competitors.

Figure 1 demonstrates this:

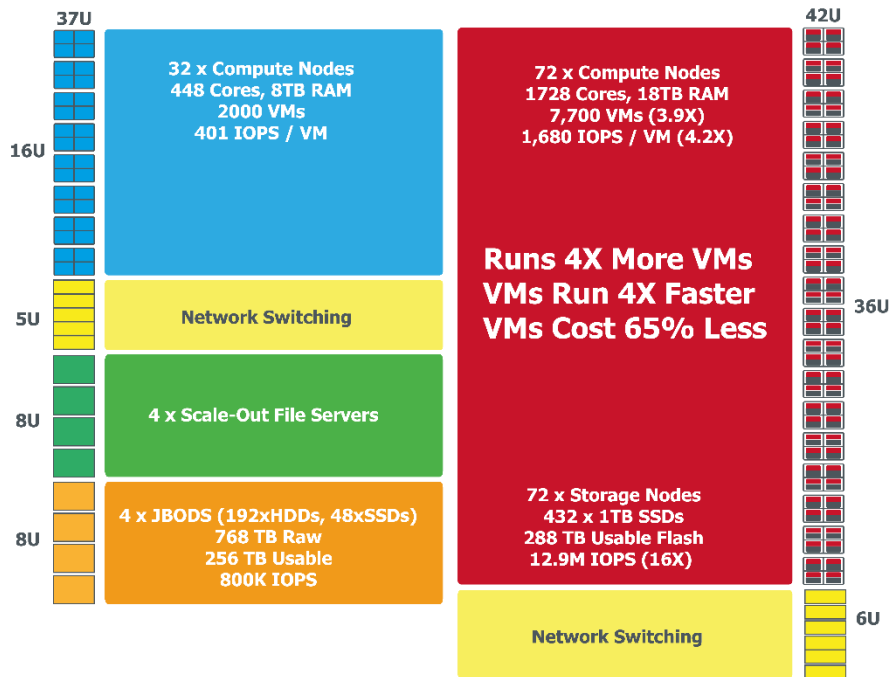


Figure 1: Microsoft CPS vs Gridstore (Gridstore 2014)

On the left, is an example of Microsoft’s recently announced Cloud Platform System (CPS). On the right, is Gridstore’s equivalent solution using the exact same Microsoft software stack.

Microsoft’s CPS delivers 401 IOPS/VM, while Gridstore, in the same footprint, delivers 1,680 IOPS/VM. Gridstore also calculates that Microsoft’s CPS costs \$825/VM, while its solutions costs 65% less at \$290/VM.

Neuralytx believes that the use of contemporary technologies, such as erasure coding, enables Gridstore’s solution to deliver higher performance, greater data storage capacity at a significantly lower cost.



Independent Scalability

Gridstore's independent scaling allows customers to scale performance and data storage capacity together independently.

This allows customers to scale only what is required for their environment. This dramatically lowers capital investment and by extension, the TCO.

Gridstore's independently scales predictably and is repeatable.

Customers have the option to add more HyperConverged Appliances if more compute performance is required or Capacity Storage Nodes if more data storage capacity is necessary.

This maximizes the agility to meet the business values needs of the customer.

Conclusion

HCI is not hype.

Neuralytix forecasts that by 2020, close to 50% of all new server deployments will be in some form of converged infrastructure, with HCI making up a large proportion of the systems shipped.

Those customers that do not move to this approach will find themselves less competitive and struggle in their respective markets.

We believe that Gridstore's solutions meets the needs of enterprise that are looking to drive the most value out of their IT investments. Specifically, Neuralytix believes that Gridstore meets the critical tests relating to predictability, repeatability and scalability, in a way that is unique, specifically Gridstore's

- Ability to deliver QoS on a per VM basis to meet service level agreements (SLAs);
- Lower TCO by as much as 50%; and
- Independent Scaling.

These features, combined in a compact HCI package makes Gridstore a solution that all growing enterprises should evaluate.

Note:

All rights for trademarks and names are property of their respective owners.

About the Author

Benjamin Woo is the founder and Managing Director of Neuralytix, Inc. He is a recognized, celebrated, provocative market visionary and thought leader. Mr. Woo frequently speaks at industry and customer events worldwide and is often quoted by leading business and technology press.

Mr. Woo has dedicated his entire career to the data and information industry. He uses this diverse set of skills, knowledge and experience to provide in-depth market insight and advice on key aspects of the IT marketplace to both vendors and buyers of technology. Mr. Woo also advises Wall Street clients and other interested stakeholders.

Mr. Woo serves on a number of advisory councils of leading storage manufacturers advising them on strategies and direction relating to the industry.

Prior to founding Neuralytix, Mr. Woo was the Program Vice President of IDC's Worldwide Storage Systems Research, where he led a team of analysts responsible for advising clients on the evolution and trends related to data storage system. While at IDC, Mr. Woo also initiated the research on Big Data.

About Neuralytix™

Neuralytix is the global leader in contemporary and relevant IT market research and consulting. We take a holistic and forward-looking approach to research, which makes us unique and the most relevant research firm in the IT industry today.

Neuralytix is the standout leader in contemporary IT market research and consulting. Our research helps enterprise end-users, vendors and the financial community make the best possible decisions as it relates to their investments in technology. It releases them from anchoring their decisions in archaic and out-of-date segmentations that only serve to hold back opportunity. Instead, our contemporary and forward-looking views sets up the framework that optimizes immediate and future gains, competitive advantage and enterprise value.

Neuralytix, Inc.

1129 Lexington Avenue, #3
New York, NY 10075
USA

Paperless Productivity
1402 Third Ave.
Suite 812
Seattle, WA 98122
Tel 1.888.838.0042
info@PaperlessProductivity.com

www.PaperlessProductivity.com

