



The HC3 Difference

The Factors that Make HC3
Stand Out from the HCI Crowd

From the Experts at
Scale Computing

Technical Whitepaper

Table of Contents

- Introduction 3
- What Makes a Real HCI Solution? 4
 - The Hypervisor..... 4
 - The Convergence 4
 - Is HC3 Real HCI?..... 4
- How is HC3 Simpler? 5
 - Deployment 5
 - Management and Monitoring 5
 - State Machines and Self-Healing 6
 - System Updates..... 6
 - Scale-Out and Refresh..... 6
 - The Little Things: Snapshots and Cloning 7
- How is HC3 Storage Different? 7
 - Automatic Configuration 7
 - SCRIBE..... 8
 - HEAT 8
 - The HC3 Storage Difference 9
- How is HC3 Different for Edge Computing? 10
 - Right-Sized Infrastructure 10
 - Total Cost of Ownership..... 11
- What Do Our Customers Say About HC3?..... 11
 - Success Stories 11
 - User Reviews 11
 - Net Promoter Score 12
- Summary 12



Introduction

Scale Computing HC3 is a hyperconverged infrastructure (HCI) solution that has been making waves across the IT industry. In 2020, Scale Computing received top honors in the CRN ARC awards in the categories of Converged/Hyperconverged Infrastructure and Edge Computing. In 2019, Scale Computing displaced some of the biggest names in the industry by winning the Desktop and Server virtualization and Converged/Hyperconverged Infrastructure categories.

	Desktop & Server Virtualization	Converged/Hyperconverged Infrastructure	Edge Computing
2020	n/a (only apply for two)	Scale Computing	Scale Computing
2019	Scale Computing	Scale Computing	n/a
2018	Dell/VMware	Dell/EMC	n/a
2017	Dell/VMware	Cisco	n/a
2016	Dell/VMware	HPE (Simplivity)	n/a
2015	Dell/VMware	Dell/EMC	n/a
2014	Microsoft	Dell	
2013	VMware	Cisco	
2012	VMware	n/a	
2011	VMware	n/a	
2010	VMware	n/a	



How could a smaller (relative to the previous winners) and emerging vendor like Scale Computing suddenly take these top spots? What are the secret ingredients in Scale Computing HC3 that make it stand out? This paper explains why so many organizations have selected us as their trusted infrastructure solution.

What Makes A Real HCI Solution?

You've probably heard of hyperconvergence or hyperconverged infrastructure and these are becoming hot new industry buzzwords but what do these terms really mean? Are vendors that say they have hyperconverged infrastructure really living up to the promises of hyperconvergence?

The Hypervisor

The term hyperconverged means the inclusion of a virtualization **hypervisor** into a complete infrastructure appliance solution that included storage, compute, and virtualization. Some have thought that hyperconverged is synonymous with terms like "ultra"converged or "super"converged but that was not the intention. The hyper in hyperconverged is the hypervisor.



What does this definition mean to be a real hyperconverged solution? Many solutions that call themselves hyperconverged rely on third-party hypervisors such as VMware or Hyper-V for virtualization. The hypervisor software, in these cases, is developed and licensed from a completely different vendor. That doesn't seem to fit the definition of hyperconvergence.

A real hyperconverged solution with a fully integrated and native hypervisor allows for tighter integration of components such as storage architecture being embedded in the hypervisor, better automation of management tasks and self-healing capabilities, and better, more seamless system update capabilities. There are many benefits to being truly hyperconverged.

One benefit of real HCI that we often hear from HCI users is independence from these third-party virtualization vendors and their software licensing. They appreciate not having separate hypervisor software licensing to pay for and one less vendor to deal with for support and maintenance. It is a complete virtualization and infrastructure solution from a single vendor.

Vendors using third-party hypervisors are doing so both to appeal to IT professionals who are already experienced with these hypervisors and because it is far easier to slap on an existing hypervisor than to create a new hypervisor that is truly integrated. Whatever their reasons they are not fulfilling the promise of hyperconvergence for their customers, and supporting third-party hypervisors leads to other inefficiencies that will be covered later in this paper.

The Convergence

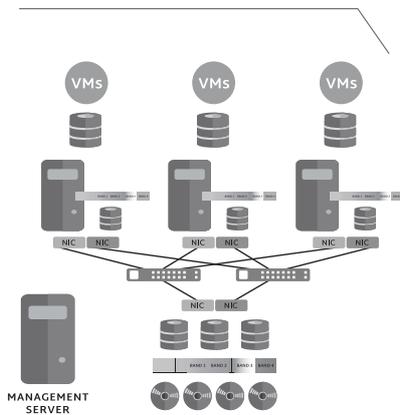
You may have more recently heard the term "disaggregated HCI" floating around. This refers to supposed "HCI" solutions where the storage is separate from the virtualization hosts. These solutions are not converged by any means. HCI, by definition, combines the server compute, the storage, and the hypervisor into a single appliance. It resembles a single virtual host server, but it is far more because of its ability to scale out quickly in clusters without requiring secondary shared storage.

Disaggregated HCI is trying to put the label of HCI on a traditional distributed infrastructure in which the servers are separate from the storage. It has more components to manage and fails both the converged definition and the full benefits of hyperconverged. This architecture increases the number of machines to buy, scale-out, and refresh. It often requires a minimum of five appliances (three storage and two compute) well above the minimum of true HCI systems. It is not HCI.

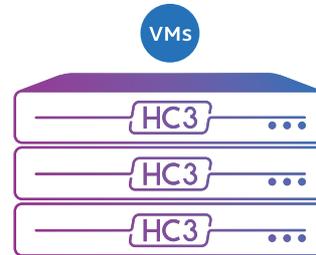
Is HC3 Real HCI?

Scale Computing HC3 is a true HCI solution with our own patented HyperCore operating system and hypervisor tightly integrated to our own SCRIBE storage architecture. The HC3 solution is delivered as an appliance integrating all of the components needed to run virtual machines and scale-out seamlessly as needed.

Virtualization The Old Way



Virtualization The HC3 Way



Not only does HC3 deliver a full virtualization infrastructure but it includes the HC3 Data Protection Suite with features like snapshots, remote replication, remote failover, failback, and recovery options for VMs or individual files. This native integration of so many infrastructure components sets HC3 apart from nearly all HCI solutions in the market.

Meeting this definition of a real HCI solution is only the start of what differentiates HC3. The real benefits are found in the simplicity, the storage efficiency, the flexibility to go from the data center to the edge, and more covered in this paper.

How is HC3 Simpler?

Simpler is not always associated with being better. IT professionals are known for wanting to have a button, knob, or switch for every possible option and sometimes those are indeed useful for certain tasks. But when it comes to everyday management tasks, fewer buttons to push and fewer consoles to monitor can save valuable time that could be better spent on other projects.

HC3 simplifies IT management using streamlined workflows and automation starting from deployment continues all the way through the lifecycle with self-healing and easy management, and also extends to scale-out and refresh. Here we will highlight how those simplifications benefit our HC3 users.

Deployment

As HC3 is designed to be a full virtualization infrastructure solution, the deployment is designed to get our users from unboxing to creating VMs as easily and quickly as possible. Aside from racking and cabling, appliance setup is as easy as assigning IP addresses, entering a license code, and then pointing a web browser to the appliance. To create a cluster, additional appliances are pointed to the IP address of the first appliance with a join command.

Unlike other HCI solutions, there is no manual storage configuration needed (more on that later). There are no hypervisors to install or third-party management consoles to setup. As HC3 management is web-based, there are no client components to install, only a web browser is required. A three-node HC3 cluster can be racked, cabled, and configured in less than an hour by an experienced engineer with live VMs running.

Management and Monitoring

With HC3, there is a single management interface to manage servers, virtualization, storage, and data protection. HCI vendors that use third-party hypervisors share management between third-party management consoles for virtualization and their own consoles for storage management.



HC3 web-based management provides live management and alerting for HC3 clusters and single appliances and also provides additional views for multi-site monitoring. For more advanced management and automation, HC3 provides a fully documented set of REST APIs accessible from the HC3 web interface.

State Machines and Self-Healing

The unique architecture of HC3 includes active state machines which monitor each appliance and are able to take action to automatically correct or at least report failure and error conditions. The benefit of owning the entire computing stack from storage to hypervisor is that it allows the HC3 architecture to have tighter integration between system components to monitor and automatically manage them more closely.

Automated, self-healing corrective actions by the state machines as well as the hypervisor high availability clustering mean that an HC3 system can continue operating despite many error and failure conditions without manual intervention. Late nights and weekends don't have to be interrupted by every system error or failure, nor does productivity have to be affected. HC3 systems are built with intelligent automation to be reliable and resilient.

System Updates

Performing firmware and software updates to systems can be one of the most important and dreaded tasks in IT. They can involve long weekends for IT professionals and taking critical workloads offline. With HC3, these system updates are completely automated and can be performed on a cluster without taking any workloads offline.

HC3 updates include firmware and software updates for all components of the system where other HCI solutions require separate updates for each component. They often require one update for firmware, a separate update for storage software, and another separate update for the hypervisor.

Scale-Out and Refresh

When it comes to scaling out a cluster, it is as simple as initial deployment, even when the new appliance being added to the cluster has a different CPU, RAM, and storage configuration. Many other HCI solutions require new appliances in the cluster to have the same or very similar configuration to the initial cluster nodes. This limits flexibility when it comes to scaling out with more of a particular needed resource.

When it is time to refresh nodes, the flexibility of HC3 to scale-out with different and newer appliances allows for the ability to gradually phase out older appliances as the cluster is scaled out with newer appliances. Rather than rip and replace, a cluster can refresh in place through scaling out and retiring older nodes over time.

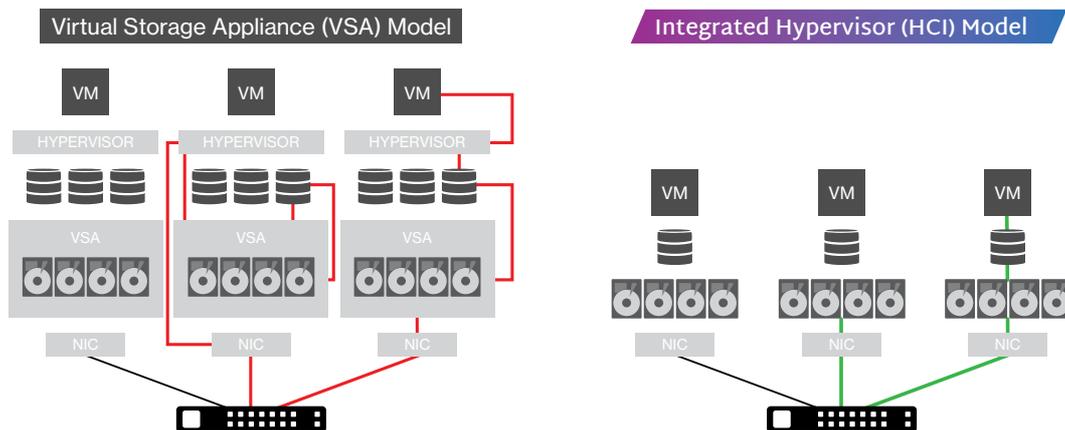
The Little Things: Snapshots and Cloning

In HC3, even some of what may be considered the basic features stand out, such as VM snapshots and VM cloning. HC3 VM snapshots are used throughout the HyperCore hypervisor for data protection features such as replication and quick VM recovery. As they are used frequently, HC3 snapshots are taken efficiently and as many as 5000 snapshots may be stored per VM.

Unlike other virtualization solutions, HC3 VM clones do not have child-parent dependence on the original VM. Clones are created thinly (sharing data blocks) with the original VM, but otherwise, the clone is completely independent, and deleting the original VM or other clones has no effect on the cloned VM.

How is HC3 Storage Different?

When it comes to data access for users, efficiency is better for productivity. Many factors affect how efficiently the data is delivered to users including networking, processing, and maybe the most important, storage architecture efficiency. Many storage architectures are complex, layering multiple storage protocols along with virtual appliances simply to present block storage to virtual machines.



By virtue of having a tight integration between hypervisor, hardware, and storage, Scale Computing has created a hypervisor-embedded storage architecture within HC3 to deliver high efficiency along with streamlined simplicity. This unique architecture provides a number of benefits.

Automatic Configuration

Other HCI vendor solutions require manual creation of storage groups from existing storage before VMs can be created. Some software-only HCI solutions require manual configuration of storage before the HCI solution can be installed and then more manual configuration based on the hypervisor requirements. These required manual management steps are generally due to those solutions supporting third party hypervisors.

With HC3, storage is embedded with the hypervisor and is automatically configured across the entire cluster. There is no manual configuration required for deployment or ongoing management. All of the storage in the cluster is combined into a single, cluster-wide storage pool for use by all nodes of the cluster. As new nodes are added to the cluster, their storage is automatically and seamlessly added to the storage pool with no manual configuration required.

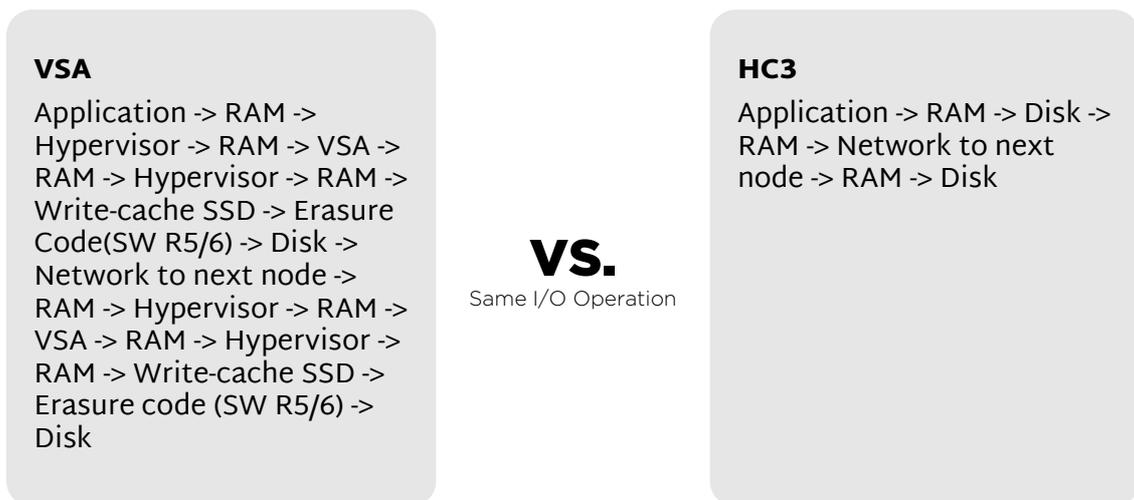
SCRIBE

The **S**cale **C**omputing **R**eliable **I**ndependent **B**lock **E**ngine is the core component of the HC3 system, combining the storage drives from each HC3 node into a single, logical, system-wide storage pool. The pooling occurs automatically with no user configuration required. Blocks are stored redundantly across the system to allow for the loss of individual drives or an entire system node.

The SCRIBE storage pool is available to all nodes of the system and presented without any file systems, protocols, or virtual appliances. SCRIBE is embedded directly in the HC3 HyperCore operating system. When a VM is created on HC3, the virtual disks provide direct block access between the virtual machines and the SCRIBE storage pool. The only file systems created in HC3 are the file systems used by the guest operating systems in the VMs to address the virtual disks.

Other storage architectures tend to emulate SAN or NAS devices that were traditionally used in virtualization. These start with a storage pool on the lowest level with a file system layered on, then presented to the hypervisor where another file system is layered on and managed by a virtual storage appliance (VSA), and finally presented to a VM where another file system is layered on. Aside from the multiple levels of protocols that must be traversed for each I/O operation, the VSAs managing the storage can consume a large portion of the RAM (often in the range of 24GB+ per cluster node) that would otherwise be used for creating more VMs.

As a comparison, the route of an I/O operation in a VSA architecture may look something like:



The example above includes an SSD cache as part of the VSA architecture. SCRIBE can also incorporate SSD but instead of using it merely as a cache, it is used as a storage tier within the SCRIBE storage pool. As a cache, the higher speed of SSD helps mask the inefficient design in VSA architecture. In HC3, as a storage tier, SSD is used for data storage, increasing both the size and the overall speed of the storage pool access. With SCRIBE, a hybrid storage architecture allows for data to be moved dynamically between SSD and HDD tiers using HyperCore Enhanced Automated Tiering (HEAT) technology.

HEAT

Data sets are most commonly characterized by only a small percentage of data being actively accessed. In a hybrid, tiered storage system, the most benefit to efficiency comes from the active data residing on SSD and the inactive data being stored on slower HDD storage. The HEAT technology in SCRIBE monitors data access and creates a dynamic mapping of active data blocks and moves these blocks to the SSD storage tier while moving inactive storage blocks off of SSD onto HDD.



This movement of data blocks between tiers is transparent to users and even administrators, happening automatically. By default, each disk is actively trying to achieve an equal level of efficiency by moving data blocks between SSD and HDD tiers. Sometimes, one disk requires an even greater level of efficiency than another disk; in these cases, HC3 administrators can give those virtual disks priority.

For each virtual disk, the relative priority of SSD utilization can be adjusted on a scale from 0-11. The default for all disks is 4 and if no disk is ever changed, no disk has priority over any other. Any disk can be dynamically adjusted in priority, up or down, to increase or decrease the SSD utilization on that disk. Priority 0 will completely bypass SSD and only utilize HDD for that disk. Priority 11 will attempt to put the data for that disk completely on SSD, or at least as possible given the available SSD on the system, taking into account all other disk priorities.

A screenshot of a dialog box titled "Update Block Device" with a close button (X) in the top right corner. The dialog contains two main settings: "Type" is set to "VIRTIO" in a dropdown menu, and "Flash Priority" is set to "11" on a horizontal slider. Below the slider, the text "These go to 11" is displayed. At the bottom of the dialog are two buttons: "Update" and "Delete".

HC3 has built-in, real-time IOPS meters for each virtual disk so the results of changing the priority on a particular disk can begin to be seen immediately. Because the virtual disks can be changed dynamically, the priority can be adjusted as much as necessary to achieve the correct balance of IOPS efficiency across virtual disks and VMs.

The HC3 Storage Difference

By design, embedding the storage system in the hypervisor allows HC3 to have extreme efficiency with absolute simplicity. With HC3, users experience the benefits of SSD as a tier, and have the benefit of clustered redundancy that absorbs the loss of failed disks or entire nodes.

This is the advantage of HC3. A storage system that is designed specifically to work with modern virtualization. None of the complexity of SAN and NAS that were designed to work with separate physical servers. Efficiency, convergence, and user experience have been simplified in one system.

How is HC3 Different for Edge Computing?

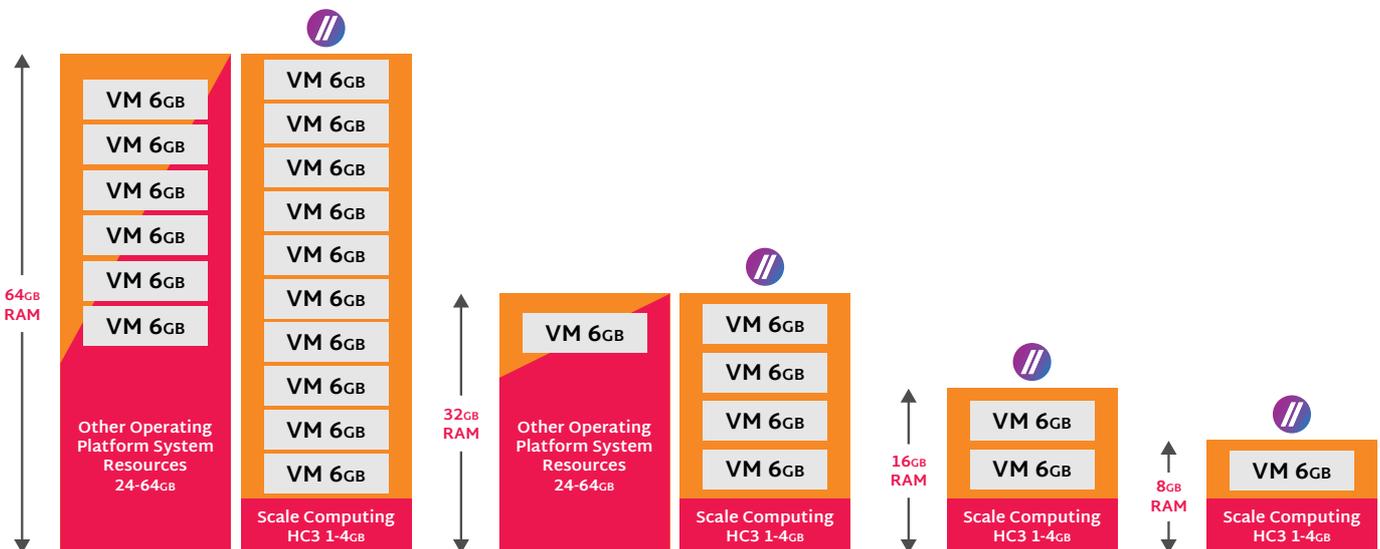
Edge computing describes a physical computing infrastructure that is intentionally located outside the four walls of the data center. The purpose of edge computing is to place applications, compute, and storage resources near to where they are needed, used, and where data is collected. Being outside the data center, the infrastructure requirements for edge computing can be much smaller, smaller even than a typical small server configuration.

Right-Sized Infrastructure

As discussed previously, the VSA-style storage architectures of other HCI solutions can consume large amounts of RAM and even CPU. These solutions simply can't scale down to smaller server/appliance form factors because of their resource consumption.

HC3 HyperCore is extremely lightweight compared to other hyperconverged virtualization stacks, mainly because of the storage architecture. HyperCore, in which the SCRIBE storage layer is embedded, only allocates about 4GB of RAM to operate. This has allowed Scale Computing to offer the HE150 HC3 Edge appliance which starts with as little as 8GB of RAM for edge computing.

Compare this to other HCI solutions that require virtual storage appliances (VSAs) which are entire VMs consuming upwards of 24GB+ of RAM before you even start creating VMs. You aren't going to find those solutions available in a small appliance like the HE150, especially where you may be targeting an appliance with 8-16GB of RAM. Even on systems with a minimum of 32GB of RAM, the HCI solution software and the VSA appliance are going to consume most of the RAM.



With this reduced resource efficiency, other HCI solutions will require larger servers to provide edge computing, impacting the total cost of ownership.

Total Cost of Ownership

In edge computing deployments, there can be hundreds or thousands of sites requiring on-premises infrastructure. Having to deploy a large, costly infrastructure stack at each site will multiply costs quickly. It will likely be cost-prohibitive to provide fault tolerance due to the unnecessary resource overhead of a larger-than-necessary deployment.

With HC3 and appliances like the HE150, a highly available, three-node appliance cluster can be deployed for the same cost as a small server and or another vendor's HCI appliance. That added high availability translates into loss of productivity or business from downtime and less time spent dealing with recovering from downtime at remote edge computing sites. We've already discussed how easy it can be to deploy HC3 systems, further reducing the costs multiplied across sites.

Consider these multipliers:

- An extra 10 minutes of management time per site x 240 sites adds up to a 40 hour work week
- An extra 4 hours of deployment time per site x 240 sites is an extra 6 months added to deployment
- An extra 24GB of RAM resource overhead per site x 240 sites is 5.7TB of resource overhead across sites
- An extra hour of downtime costing \$1000 of lost business per site x 240 sites adds up to a loss of \$240,000.

Considering these factors, no other HCI solution can really compare to HC3 in being able to provide an easy to manage, right-sized, and cost-effective solution for edge computing.

What Do Our Customers Say About HC3?

Going through the buyer's journey, there is plenty of marketing and technical content to consume but there are other factors to consider, and one that is considered most often is how other customers view the vendor and their products. Like most vendors, we produce case studies and success stories, and we have HC3 users across nearly every industry willing to be customer references, but independent third party review sources and ratings are also just as important when considering HC3.

Success Stories

Without going into too much detail, there are customer success stories published for use cases across nearly every industry and around the world. The purpose of these success stories is to highlight the challenges being faced by the customer, why the customer put trust in HC3 to meet their needs, and how HC3 solved their challenges.

Please see the full list of our [published customer success stories on our website](#).

User Reviews

One of the most important factors in choosing an infrastructure solution for those we talk to is reviews from peers. They want to know about the experiences of our HC3 users, good or bad. Where our prospective customers often look for reviews is not from our own resources, but from external sources like [Gartner Peer Insights](#), [Spiceworks](#), [TrustRadius](#), and others.

These sources validate the reviews independently and we encourage our HC3 users to share their experiences. We know that based on our customer success with HC3, that our reviews will exceed the ratings of other HCI solutions, and they do. While we could pull in our own hand-picked customer quotes and reviews for you to see here, we instead encourage you to use the links above to look through the reviews with all of the pros and cons listed. You will find that HC3 really is highly rated and reviewed compared to competing solutions.

Net Promoter Score

Net Promoter Score (NPS) is a rating based on how likely a customer will be to recommend your company or products to someone else. It is a rating used across nearly all industries to determine levels of customer satisfaction. The scores can range from -100 to 100. You can find a good description of how these scores work and what they mean in this [blog post from Jamie Yan on questionpro.com](#).

Our NPS at Scale Computing is constantly being updated based on customer surveys we conduct when in contact with our customers through our support team. Despite the constant fluctuation, our average NPS has remained consistent for years in the 80+ range, usually in the high 80s or low 90s. That is well above the world-class standard for NPS. Our customers will tell you that this is a combination of both our excellent customer support and our product quality.

Summary

Hyperconverged Infrastructure is rapidly replacing traditional virtualization infrastructure for primary data center computing, edge computing, and virtual desktop infrastructure (VDI). As HCI continues to grow and evolve, it becomes more clear who the real HCI players are and who the HCI pretenders are. HC3 not only leads HCI and virtualization solutions in reviews and ratings, but our unique patented architectures and technologies are changing the way consumers of IT infrastructure think about how HCI can be flexible enough to meet almost any infrastructure challenge.

If you are ready to learn more about HC3 for data center and edge computing solutions, our experts are ready to answer your questions. We can provide demonstrations of our current solutions and even prototypes of new technologies yet to hit the market. We want to hear from you not just to answer your questions but to hear your unique computing challenges and insights into how we can make HC3 an even better solution.

For more information on how to get started with your HC3 Edge solution or if you have additional questions, contact Scale Computing at info@scalecomputing.com or call 877-722-5359.

Corporate Headquarters
525 S. Meridian Street - 3E
Indianapolis, IN 46225
P. +1 317-856-9959
scalecomputing.com

EMEA B.V.
Europalaan 28-D
5232BC Den Bosch
The Netherlands
+1 877-722-5359

