

ESG Lab Review

Pivot3 Acuity: A High-performance Hyperconverged Platform with Advanced QoS

Date: December 2017 Author: Kerry Dolan, Senior IT Validation Analyst

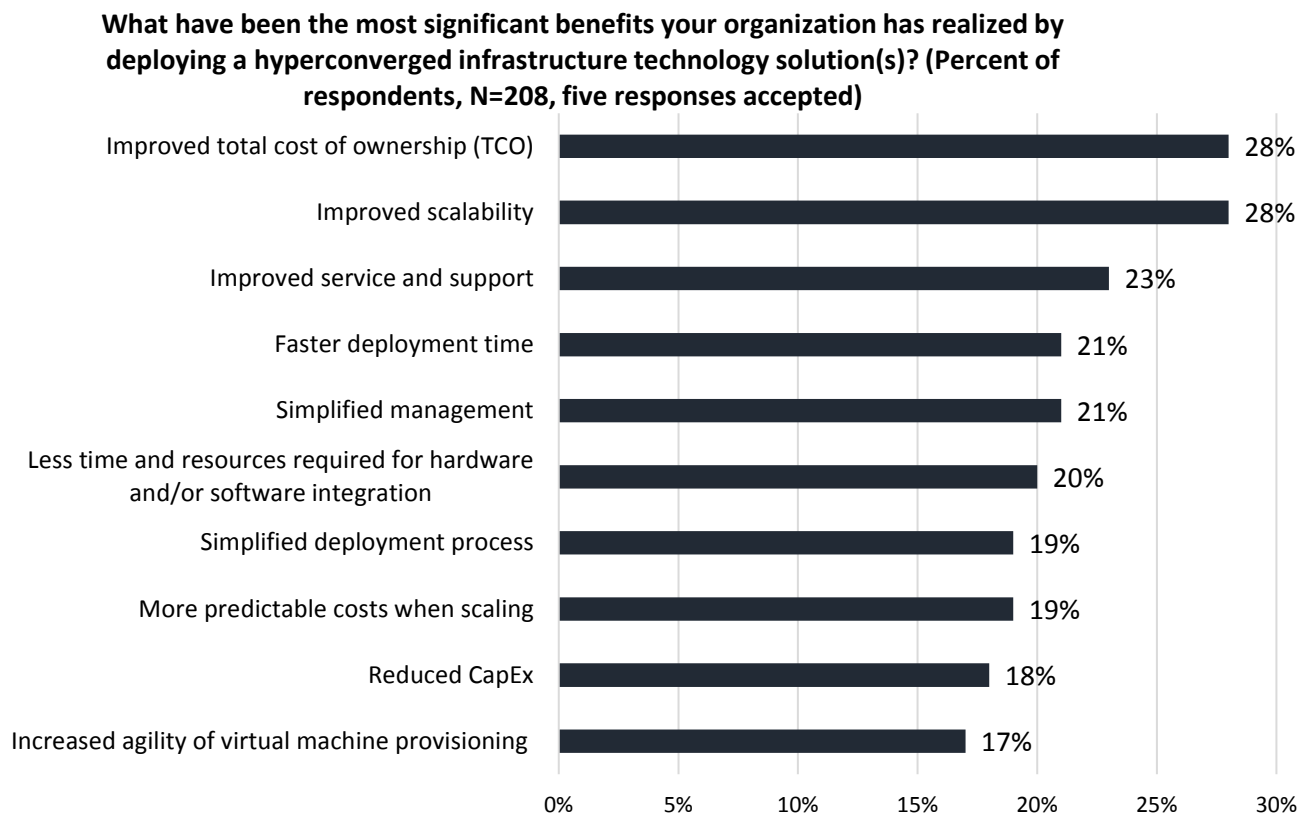
Abstract

If organizations can maintain predictable high performance that is prioritized by workload, they can deploy mixed applications on hyperconverged infrastructure (HCI) deployments to increase efficiency, reduce costs, and simplify management. This ESG Lab Review documents validation of Pivot3 Acuity HCI, with NVMe PCIe flash-enabled high-performance and policy-based quality of service (QoS) to maintain predictable, workload-prioritized high performance to support mixed application workloads.

Background

Organizations struggle to achieve IT agility with individually managed infrastructure silos. Many are turning to HCI, which enables them to centrally manage a single unit of software-defined compute, network, and storage that is flexible, scalable, and easy to deploy. And the rewards are plenty: ESG survey respondents report achieving benefits that include better TCO, scalability with predictable cost, faster and simpler deployments, simpler management, reduced CapEx, and more agile virtual machine (VM) provisioning (see Figure 1).¹ It sounds too good to be true!

Figure 1. Top Ten Benefits Realized by Deploying HCI



Source: Enterprise Strategy Group, 2017

¹ Source: ESG Master Survey Results, [Converged and Hyperconverged Infrastructure Trends](#), October 2017.

This ESG Lab Review was commissioned by Pivot3 and is distributed under license from ESG.

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

And in many cases, it *is* too good to be true. To get to market fast, some HCI solutions traded performance for simplicity, leaving them inadequate for mission-critical workloads. In addition, in response to poor mixed workload performance due to resource contention, many organizations limit HCI deployments to a few targeted applications. So while they are converging infrastructure, they are not achieving the greater efficiency gains of converging applications.

The Solution: Pivot3 Acuity

Pivot3 Acuity is an HCI platform that combines a high performance NVMe data path and multiple storage tiers with advanced, policy-based QoS to enable the consolidation of multiple, mixed application workloads. Available in all-flash and hybrid configurations, Acuity leverages NVMe PCIe flash, RAM, SSDs, and HDDs to automatically deliver the right performance for workloads based on their business priority. NVMe is a standardized host controller interface that is designed to maximize flash performance, eliminating the inefficiency and latency of legacy SCSI interfaces. Acuity nodes are clustered and each node’s capacity, IOPS, bandwidth, and cache are aggregated and available to any VM in the cluster.


Acuity performance QoS goes beyond simple capping of IOPS or throughput. Acuity QoS capabilities include:

- Five pre-set QoS policies that define minimum IOPS, throughput, and not-to-exceed response times
- I/O prioritization based on service-level designation (Mission-Critical, Business-Critical, Non-Critical)
- Real-time data placement on multiple storage tiers based on service-level designation
- Policies can be quickly changed in real time or scheduled to change to support an application’s high activity periods
- QoS data protection (snapshots, clones, and replication prioritized by service levels)

Other Acuity features include patented erasure coding, inline data reduction, and vSphere/vCenter integration.

Customers can start with three Acuity nodes and scale the system with additional nodes to meet their performance and capacity needs. Acuity nodes include dual 18-core Intel E5 CPUs, up to 1.5TB RAM, up to 30TB SSD, up to 48TB HDD, and up to 3.2TB of NVMe flash capacity.

Figure 2. Pivot3 Acuity HCI



Source: Enterprise Strategy Group, 2017

ESG Lab Tested

ESG Lab audited testing that used industry-standard tools and methodologies to validate performance across multiple workloads, as well as predictable performance using policy-based QoS. Testing included Pivot3 Acuity all-flash and hybrid configurations.

Performance

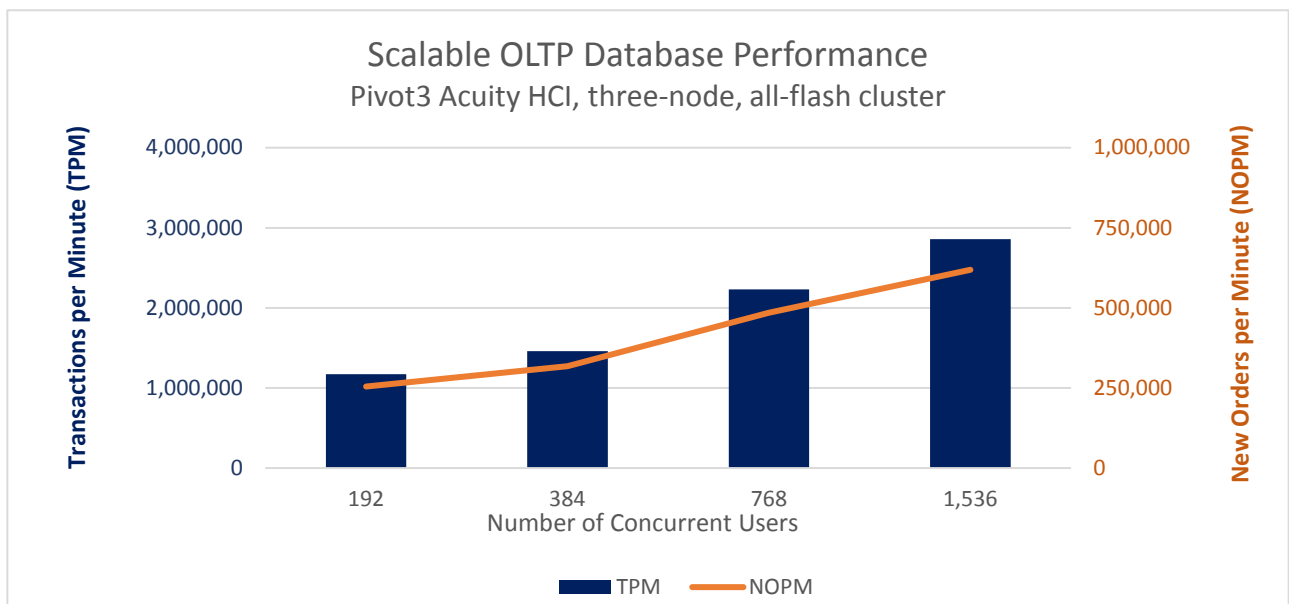
Database

ESG Lab began by auditing database performance testing that used HammerDB. This open source database benchmarking tool comes prepackaged with an online transaction processing (OLTP) workload. The benchmark is centered around the five principal activities (transactions) of an order-entry environment: Entering orders, delivering orders, recording payments,

checking the status of orders, and monitoring warehouse stock levels. HammerDB was configured for 5,000 data warehouses; as the number of users scaled exponentially (from 1-512), we captured transactions-per-minute and new-orders-per-minute results.

The Acuity system under test was a three-node, all-flash cluster comprised of two Acuity X5-6500 Accelerator nodes, each node with 512GB RAM, 1.6TB NVMe PCIe flash, and 6.4TB SSD, and one Acuity X5-6000 Standard node with 512GB RAM and 6.4TB SSD. Users were spread across three SQL Server 2014 VMs, one VM per node, configured with 32 CPUs and 384GB RAM. Figure 3 shows that with 512 users on all three VMs, for a total of 1,536 users, the Acuity configuration supported an impressive 2.8M transactions per minute, and just under 619K new orders per minute.

Figure 3. HammerDB Database Performance



Source: Enterprise Strategy Group, 2017

Virtual Desktop Infrastructure (VDI)

Next, ESG Lab looked at VDI, a common HCI application. VDI workloads are often constrained by storage I/O, so organizations typically over-provision to meet VDI performance needs. With Pivot3 Acuity, NVMe PCIe flash increases performance and desktop density, eliminating the need to over-provision storage performance. Resource aggregation makes Acuity scale in a modular fashion with predictable performance. This is valuable for VDI, where organizations typically start small and grow incrementally. From an infrastructure standpoint, administrators simply add Acuity nodes and the new resources are added to the pool.

ESG Lab validated internal Pivot3 results of Login VSI testing to evaluate Acuity’s ability to scale the number of virtual desktops while maintaining a performance threshold. Login VSI mimics real-world users performing typical tasks. We validated all-flash and hybrid deployments, with the systems under test in isolated networks. Both test beds used a three-node cluster for management and the Login VSI test harness. The Acuity test beds included:

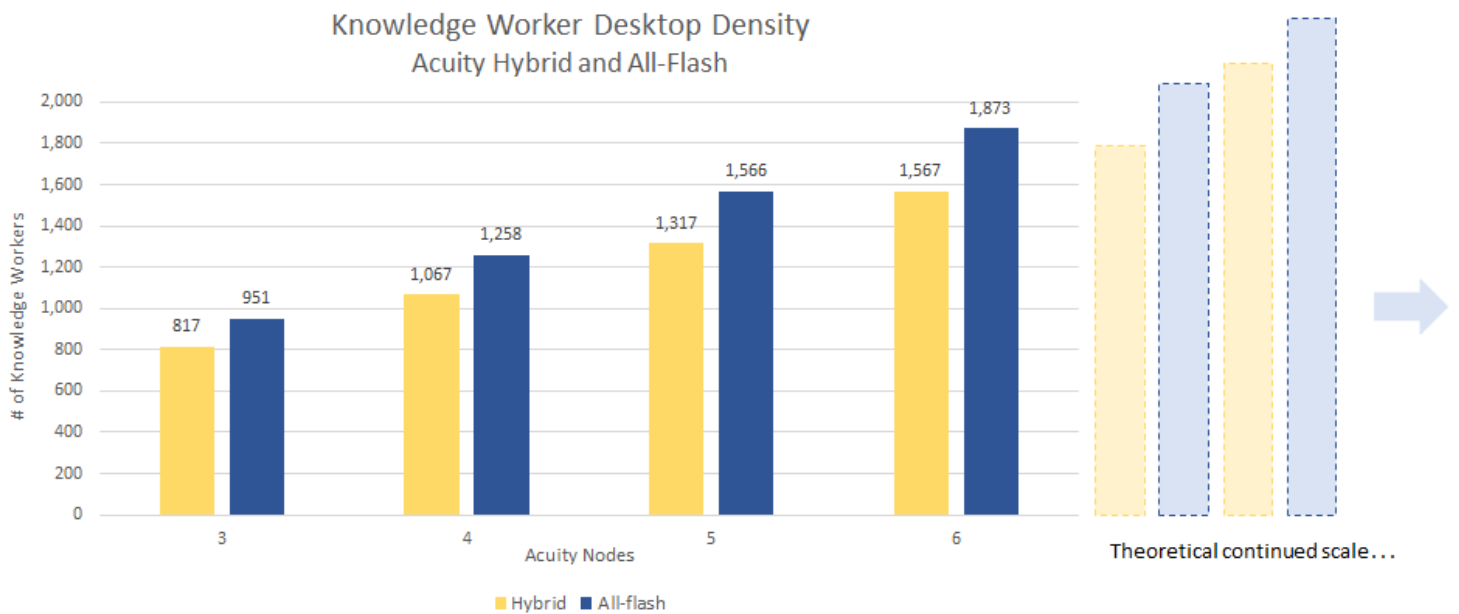
- All flash: 6-node cluster containing a total of 216 Intel E5 2695 CPU cores, 4.6TB RAM, 3.2TB NVMe flash, 76.8TB SATA SSD

- Hybrid: 6-node cluster containing a total of 216 Intel E5 2695 CPU cores, 4.6TB RAM, 3.2TB NVMe flash, 72TB SATA HDD
- Software: Acuity version 2.1; VMware ESXi 6.0, Horizon view 7.0.2; Login VSI 4.1.25.

As Login VSI runs, it boots up VMs and mimics activities of different types of workers;² VMs are considered successful once they launch, boot the operating system, begin typical tasks (e.g., opening documents, printing, browsing, or playing a video), and achieve a minimum performance threshold that is calculated early in the test. Testing was conducted in Login VSI Benchmark mode, in which test settings cannot be adjusted to improve performance.

Figure 4 shows results for knowledge worker desktops, a taxing workload. Pivot3 Acuity delivered linear scalability, adding 250 knowledge workers as each node was added to the cluster. With six nodes, it supported 1,567 knowledge workers for the Acuity hybrid configuration and 1,873 for the flash. A close examination of the hybrid configuration showed that while response time increased with the number of knowledge workers, average response time was 704 ms, well under the 1,704 ms threshold for the test. ESG Lab also validated CPU readiness and usage; while CPU usage was at times close to peak (as would be expected), no VMs waited for CPU in order to function.

Figure 4. Desktop Density Scaling



Source: Enterprise Strategy Group, 2017

In all cases, desktop density increased in a linear fashion while maintaining response time. Table 1 shows the numbers of each type of worker desktops supported for Acuity hybrid and flash configurations as nodes increased (the knowledge worker data is highlighted).

² Desktop configurations for tested workers, all using Office 2010: Task worker: Windows 7 32-bit, 1 vCPU, 1GB RAM; Office worker: Windows 7 64-bit, 1 vCPU, 2GB RAM; Knowledge worker: Windows 7 64-bit, 2 vCPU, 2GB RAM; Power user: Windows 7 64-bit, 3 vCPU, 2GB RAM.

Table 1. Desktop Density Scalability


	3 Nodes		4 Nodes		5 Nodes		6 Nodes	
	Hybrid	Flash	Hybrid	Flash	Hybrid	Flash	Hybrid	Flash
Task worker	1,162	1,261	1,495	1,614	1,829	1,967	2,162	2,320
Office worker	821	963	1,048	1,271	1,274	1,579	1,501	1,887
Knowledge worker	817	951	1,067	1,258	1,317	1,566	1,567	1,873
Power user	749	831	884	1,080	1,020	1,329	1,155	1,578

Source: Enterprise Strategy Group, 2017

Performance Characterization

Next, ESG Lab audited internal Pivot3 testing with IOMeter, an open source I/O generation tool for measuring storage performance capabilities. The system under test was the same three-node, all-flash Acuity cluster as was used for database testing. The test bed leveraged four Windows Server 2012 R2 64-bit VMs built on VMware ESXi 6.0; each VM included two vCPUs and 8GB memory.

Figure 5. Acuity Max I/O Performance Results



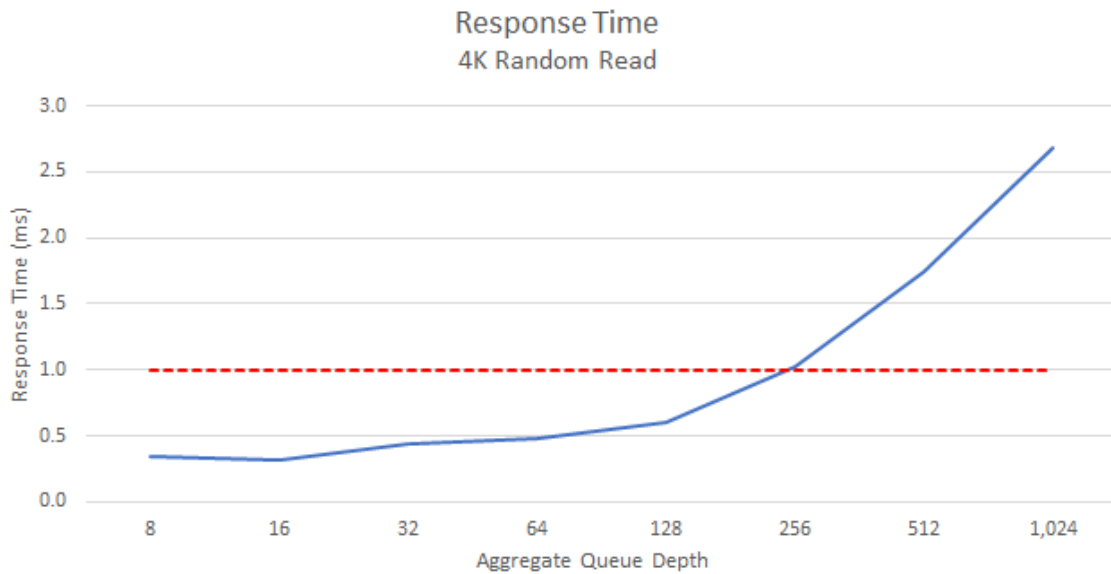
Source: Enterprise Strategy Group, 2017

Test runs were conducted with large- and small-block I/O ranging from 4KB to 256KB, random and sequential, and reads and writes. Each workload type was run with exponentially increasing per-volume queue depths from 1-128, across eight volumes (one IOMeter worker per volume), resulting in aggregate system queue depths from 8 – 1,024. Volumes were assigned to Policy 1, designed to deliver mission-critical service levels. Figure 5 displays a sample of results that represent typical workloads: small-block read/write workloads such as online databases and email, and large-block sequential workloads such as video streaming and backup. The Acuity platform delivered high read and

write IOPS for small-block workloads (top), as well as high throughput for large-block sequential reads and writes (bottom).

Next, ESG Lab looked at response time (latency) for these workloads. High queue depth indicates a system handling multiple hosts and workloads; the ability to maintain low latency at high queue depth demonstrates that the system can support multiple concurrent applications without compromising performance. Figure 6 shows response time charted for the 4KB random read workload typical of email and online databases. A good threshold response time for this workload is 1 ms, and the Acuity cluster stayed below that until an aggregate queue depth of 256 (eight workers, each at queue depth 32), an excellent result. Even at very high 1,024 queue depth (eight workers, each at queue depth 128), the Acuity cluster stayed under 3 ms. It should be noted that end-users would begin to notice delays in application responsiveness at about 20 ms.

Figure 6. Response Time as Queue Depth Scales



Source: Enterprise Strategy Group, 2017

i Why This Matters

While HCI platforms are appealing for workload consolidation, many organizations struggle to achieve the levels of performance on HCI that today’s mission-critical workloads demand. They resort to leaving HCI for applications with low performance needs, and maintain inefficient silos by keeping only single applications on HCI to avoid resource contention that drags down performance.

ESG Lab validated that the Pivot3 Acuity architecture, with aggregated NVMe data path, RAM, SSD, and HDD resources governed by its advanced QoS engine, provides the high performance that mission-critical, latency-sensitive workloads demand. This enables organizations to deploy mission-critical databases—for accounting, trading platforms, order-entry, etc.—on Acuity HCI, as well as support more VMs and VDI users per node. This increased density means that IT can eliminate performance-based over-provisioning, gain efficiency, and reduce TCO. Acuity can support consolidation of multiple applications with predictable performance, scalability, and lower cost.

Advanced QoS

The primary obstacle to consolidating multiple workloads on hyperconverged platforms is resource contention. If you place your mission-critical order database in the same HCI cluster as non-critical marketing tasks, what happens when resources are constrained? Which workload wins out? In most deployments, resources are distributed without regard to the business priority of the applications, so mission-critical workloads may become I/O-starved while non-critical workloads consume resources. Consequently, most organizations keep limited applications on HCI nodes. This results in wasted capacity and extra management, driving up both CapEx and OpEx.

While some HCI solutions provide QoS functionality, they are limited to a simple capping of IOPS and throughput: Give *workload A* 70% of resources and *workload B* 30%. But what happens when *workload A* doesn’t consume all of its allotted

I/O, while *workload B* needs more performance resources? *Workload B* would be starved of any unused performance and suffers unnecessarily.

Pivot3 Acuity’s Advanced QoS solves this challenge by setting performance minimums. Using the prior example, *workload B* would get available performance resources until *workload A* requires them again. This intelligent resource prioritization is automated, and enables the successful consolidation of multiple mixed workloads. QoS can be applied not only to VMs, but also to individual volumes within VMs. Data is automatically placed in the appropriate storage tier: NVMe, SSD, or HDD. Administrators can prioritize workloads by assigning the appropriate policy and service level, change priorities in real time, and schedule automatic policy changes to support increased resource needs such as VDI boot storms, quarterly reporting, etc.

ESG Lab audited QoS demonstrations that were performed on the three-node, all-flash X5 6500/6000 Acuity cluster. Five VMs were configured to run mixed Iometer workloads: VMs 1, 2, and 5 ran 4KB random read workloads, while VMs 3 and 4 ran 16KB mixed database workloads. Volumes were assigned QoS policies; nine volumes were designated Mission-Critical, Policy 1, with performance minimums of 125K IOPS, 1,000 MB/s throughput, maximum 1 ms latency, read-warming at 1 hit per MB, and read-ahead enabled. Other volumes were assigned Policies 2 and 3 (Business-Critical) and 4 and 5 (Non-Critical). Figure 7 shows QoS policy configurations with assigned volumes; Policies 2-5 were assigned higher latency, fewer IOPS, less throughput, and higher numbers of hits to trigger read-warming, as well as enabling or disabling read-ahead.

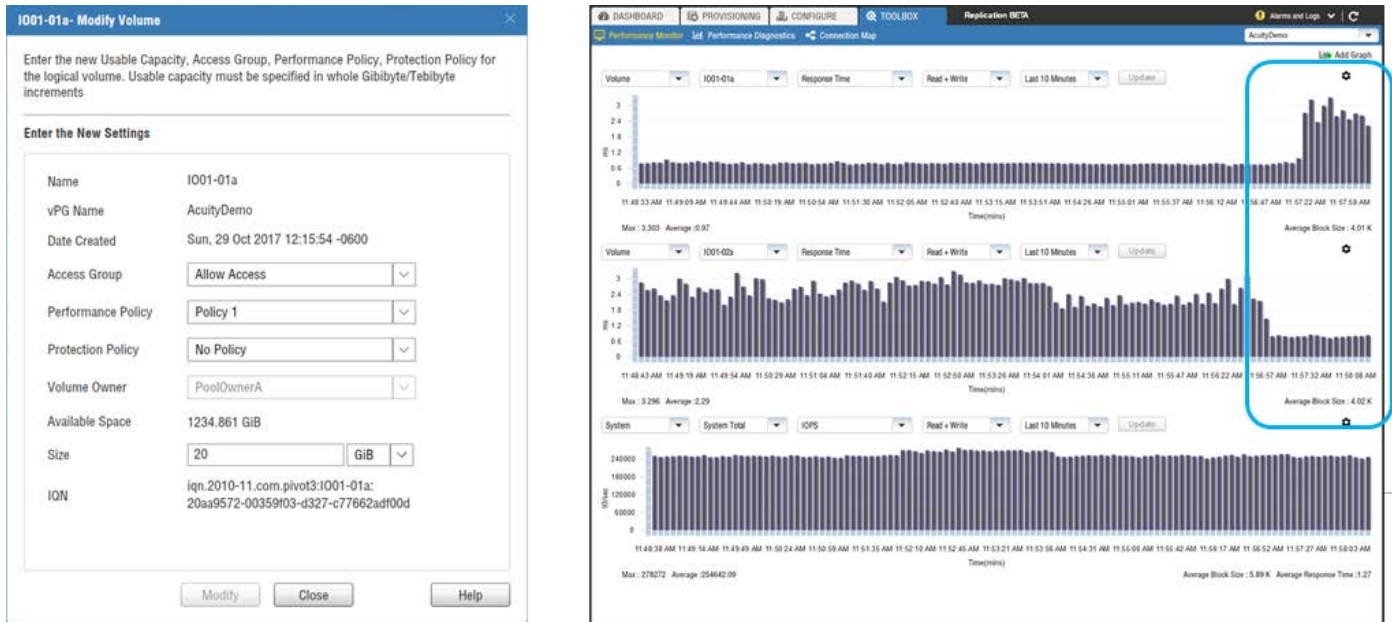
Figure 7. Pivot3 Acuity QoS Policies

Name	Service Level	Volumes	Latency Max Target (ms)	IOPS Min Target (IO/s)	Throughput Min Target (MB/s)	Read-Warm	Read-Ahead	Write-Cache
Policy 1	Mission-Critical	9	1	125000	1000	1 hit per MB	Enabled	Enabled
Policy 2	Business-Critical	0	3	75000	500	4 hits per MB	Enabled	Enabled
Policy 3	Business-Critical	3	10	50000	250	16 hits per MB	Enabled	Enabled
Policy 4	Non-Critical	0	20	25000	100	Disabled	Disabled	Enabled
Policy 5	Non-Critical	22	40	10000	50	Disabled	Disabled	Enabled

Source: Enterprise Strategy Group, 2017

Next, we changed the policies of two volumes in real time and watched Acuity shift resources. Volume IO01-01a was configured as Mission-Critical, Policy 1, and Volume IO01-02a as Non-Critical, Policy 5. We selected the two volumes, clicked *Modify Volume*, and swapped their policies. Immediately, the first volume’s response time went from less than 1ms to 3 ms; the second volume went from 3 ms to less than 1 ms, while the total system resources remained the same. Figure 8 shows the ease of modifying a volume policy on the left, and the response time changes for the two volumes, as well as IOPS for the total system remaining the same, on the right.

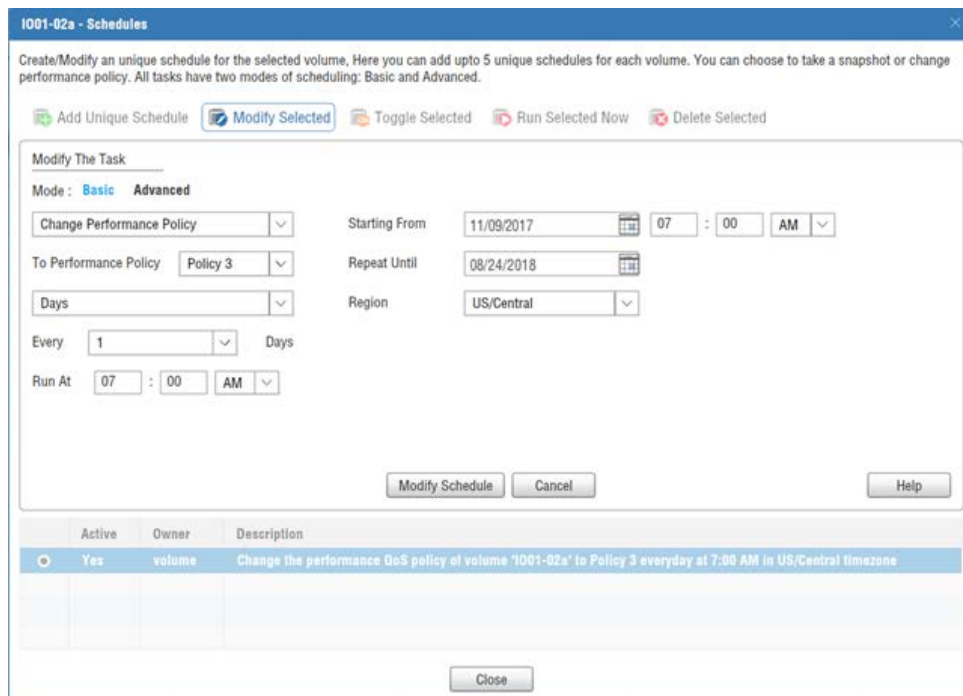
Figure 8. Real-time QoS Policy Change



Source: Enterprise Strategy Group, 2017

VDI deployments are another example of a workload that benefits from QoS. Administrators must manually optimize performance as the VDI load changes. With Acuity, they can manage and schedule resources with QoS to improve performance while reducing administrative complexity. For example, different classes of VDI users may have prioritized resources, or administrators might schedule resource changes to handle VDI boot storms without I/O contention. ESG Lab easily created a schedule for VDI boot storms that would move VDI workloads up from Policy 4 to Policy 3 at 7:00 am daily to accommodate users coming in and starting up their workstations (see Figure 9). Administrators can then create another policy when their users are likely to be booted in order to return VDI workloads to a lower policy.

Figure 9. Scheduling a QoS Policy Change



Source: Enterprise Strategy Group, 2017

Why This Matters

Unless they can guarantee performance for high-priority workloads, IT administrators have little choice but to silo applications. That leaves organizations that cannot afford performance degradation during resource contention unable to take advantage of HCI's efficiency.

ESG Lab validated that with Pivot3 Acuity's Advanced QoS, it is simple to prioritize applications with policy-based performance. Workloads are guaranteed minimum performance levels, with lower priority applications giving up resources in times of contention. This enables organizations to gain the cost, scalability, management, and efficiency benefits of HCI, mixing applications of varying priorities, while maintaining predictable performance. They can mix applications and types of virtual desktops, shift resources in real time, and schedule policy changes to support times of high activity.

The Bigger Truth

Early HCI solutions were designed to maximize simplicity of deployment and management. But they often lacked the performance levels required by today's mission-critical workloads, and organizations risked achieving SLAs by combining multiple workloads on HCI nodes.

Still, HCI deployments are gaining in popularity: According to ESG research, the percentage of our respondents currently using HCI more than doubled in the last two years, from 15% in 2015 to 39% in 2017.³ That is because HCI offers tangible benefits: By consolidating infrastructure into software-defined, centrally managed modules instead of compute, network, and storage silos, organizations gain efficiency in both footprint and management, reducing costs and complexity.

ESG Lab validated that Pivot3 Acuity HCI, with its multi-tiered NVMe flash data path, SSD, and HDD architecture, can provide the high performance that mission-critical, latency-sensitive applications demand, while retaining HCI's efficiency, scalability, and ease of management. But even more important, Acuity's policy-based, automated QoS ensures the right resources for every workload without IT intervention. This intelligent QoS guarantees performance based on workload priority, so organizations can place multiple applications—even those needing high performance—on the same HCI platform without worrying about resource contention impacting performance of critical applications.

ESG Lab was impressed with the performance and QoS capabilities of Pivot3 Acuity that enable not just infrastructure consolidation, but application consolidation. Organizations looking for a way to reduce costs and complexity would be wise to evaluate Pivot3 Acuity.

³ Source: ESG Master Survey Results, [Converged and Hyperconverged Infrastructure Trends](#), October 2017.

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. 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.

The goal of ESG Validation reports is to educate IT professionals about information technology solutions for companies of all types and sizes. ESG Validation reports are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objectives are to explore some of the more valuable features and functions of IT solutions, show how they can be used to solve real customer problems, and identify any areas needing improvement. The ESG Validation Team's expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments.