

# A Comparison of Pivot3 Datacenter Series and Dell EMC VxRail for Next Generation Datacenter Requirements

by DCIG Lead Analyst, Storage, Ken Clipperton

## PRODUCTS

### Pivot3 Datacenter Series Hyperconverged Infrastructure

URL ► <https://pivot3.com/product/datacenter/>

221 W 6th St #750  
Austin, TX 78701  
(512) 807-2666

### Dell EMC VxRail

URL ► <https://www.dell.com/en-us/converged-infrastructure/vxrail/>

176 South Street  
Hopkinton MA 01748  
(508) 435-1000

## Hyperconvergence is Reshaping the Enterprise Data Center

Virtualization largely shaped the enterprise data center landscape for the past ten years. Hyperconverged infrastructure (HCI) is beginning to have the same type of impact, re-shaping the enterprise data center to fully capitalize on the benefits that virtualizing the infrastructure affords them.

DCIG defines hyperconverged infrastructure (HCI) as a pre-integrated solution that virtualizes compute, storage and data protection functions along with a hypervisor and scale-out cluster management software. HCI vendors may offer their solutions as turn-key appliances, installable software or as an instance running on public cloud infrastructure. The most common physical instantiation of—and unit of scaling for—hyperconverged infrastructure is a 1U or 2U rack-mountable appliance containing 1–4 cluster nodes.

## HCI Requirements for Next-Generation Datacenter Adoption

The success of initial HCI deployments in reducing complexity, speeding time to deployment, and lower costs compared to traditional architectures has opened the door to an expanded role in the enterprise data center. Indeed, HCI is rapidly becoming the core technology of the next-generation enterprise data center. In order to succeed in the core, these HCI solutions must meet a demanding set of expectations.

Enterprises require core data center technologies to deliver:

1. Simplified management, including at scale
2. Workload consolidation, including mission-critical
3. High levels of Availability, Reliability and Security
4. Flexible, agile scalability

This report looks at how the Pivot3 Data Center Series and Dell EMC VxRail compare against these next-generation criteria.

## Industry Standard Hardware is Foundational to HCI Platforms

The Pivot3 Data Center Series and Dell EMC VxRail use similar hardware: Intel Xeon Scalable Processors, large amounts of RAM, NVMe, and scalable high-performance storage capacity. Thus, it is not the hardware but the respective hyperconverged software platforms that deliver differentiation. (Figure 1.)

Both solutions are built on enterprise-proven VMware virtualization, but use different hyperconverged software. VxRail uses multiple Dell EMC and VMware software packages. Pivot3 uses its own Acuity hyperconverged software platform.

Pivot3 Acuity's software was designed first as a software-defined storage platform and then was extended with a state-of-the-art software architecture developed to leverage multiple tiers of storage, irrespective of

**FIGURE 1: Pivot3 Data Center Series and Dell EMC VxRail Appliances**  
Per Node Summary Comparison

	Pivot3 2U	VxRail 2U	Pivot3 1U	VxRail 1U
Nodes per Chassis	1	1 or 4	1	1
CPU Sockets	Dual Socket	Single or Dual Socket	Dual Socket	Single or Dual Socket
RAM	256-1536 GB	64-3072 GB	192-768 GB	64-3072 GB
CPU	Intel Xeon Scalable	Intel Xeon Scalable	Intel Xeon Scalable	Intel Xeon Scalable
GPU Option	✔	✔	●	●
NVMe Flash	1.92 - 4 TB	400 GB - 1.6 TB	960 GB - 1.6 TB	400 GB - 1.6 TB
All-flash Capacity (Raw)	15.3 - 61.4 TB	1.92 - 76.8 TB	7.6 - 30.7 TB	1.92 - 30.7 TB
HDD Capacity (Raw)	12 - 144 TB	1.2-48 TB	8 - 16 TB	1.2 - 19.2 TB
Connectivity	10 GbE	1, 10, 25 GbE	10 GbE	1, 10, 25 GbE
Max Cluster Size	16 nodes, Unlimited Domain Size	64 nodes	16 nodes, Unlimited Domain Size	64 nodes

storage technology, and tightly integrated with PCIe and then NVMe flash. These technologies enable enterprises to confidently consolidate more workloads in less infrastructure, and at less cost and complexity than VxRail. These Pivot3 distinctives include:

- Simplified policy-based management that reduces OPEX
- Advanced yet simple performance quality of service (QoS) that reduces CAPEX by consolidating more workloads on less hardware while ensuring SLAs are met for mission-critical applications
- Patented erasure coding that creates highly available, reliable and efficient storage

The capabilities above are all part of Pivot3's "Intelligence Engine" which provides advanced software capabilities to simplify and automate performance, data protection, security, monitoring and analytics.

## Simplified Policy-based Management Reduces OPEX

Pivot3's automated policy-based approach to infrastructure management—including performance, data protection and security—is simple to implement yet powerful in operation. Simply assign one of five policies to a volume and the system dynamically optimizes the infrastructure to meet the pre-defined targets of every running workload. Pivot3's Intelligence Engine limits the performance of a less-important application only when it is contending for resources with a higher priority workload. (Figure 2.)

VxRail vSAN storage does not share any of these SLA management qualities.

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

**FIGURE 3: Quality of Service Summary Comparison**

Pivot3	VxRail
Rich policy-based QoS targets for latency, IOPS and throughput	Assign fixed IOPS limit to a volume
NEVER artificially limits application performance	Limits application performance even when there is no contention
Guarantees performance of high-priority workloads by throttling lower priority workloads when resource contention occurs	Does not know about priority. Lower priority workloads served up to the IOPS limit at the expense of high-priority workloads.
QoS intelligence also applies to snapshots and replication	

## Workload Consolidation through Advanced Quality of Service (QoS) Technology

Pivot3's Intelligence Engine provides the most advanced, yet simple to use QoS that our analysts have evaluated. It prioritizes workloads in real-time according to business value. This reduces CAPEX by enabling businesses to consolidate more workloads on less hardware. (Figure 3.)

In contrast, VxRail QoS is rudimentary. Administrators can assign fixed IOPS limits per volume, but that is all. These limits throttle workloads even when there is no resource contention, yet compete up to those limits with more important workloads.

### KEY QUESTIONS TO ASK:

- Does the vendor recommend dedicating hardware to support critical applications? How much could you save by eliminating that dedicated infrastructure?
- For production environments, what are the recommended utilization percentages for CPU? DRAM? raw storage capacity? usable storage capacity?
- How much—or what percentage—of the overall cost of the proposed HCI solution is associated with overprovisioning of resources?
- Does the QoS technology artificially throttle or limit applications even when there is no resource contention?
- Does the QoS implementation enable the system admin to instruct the HCI system on which applications are higher priority than others and then the system acts on those priorities in real-time to meet application SLAs?

**FIGURE 2: SLA Management Simplicity Summary Comparison**

	Pivot3	VxRail
<b>Automated policy-based SLA management of Performance, Data Protection and Security</b>	✓	●
<b>Intelligent algorithms ensure that the customer's business priorities are honored</b>	✓	●
<b>QoS policy changes can be scheduled to match time-based or calendar-based changes in workload priority.</b>	✓	●

### KEY QUESTIONS TO ASK:

- How much time does your IT staff spend attempting to address SLA-related problems?
- What costs would failure to meet SLAs impose on the company?

## Patented Erasure Coding Adds High Availability and Reliability, Subtracts CAPEX

Pivot3’s patented erasure coding and intelligent data layout yields highly reliable and available storage with a storage efficiency of up to 82% usable/raw before applying any compression. Pivot3 requires much less hardware than VxRail to deliver a required level of availability and recovers from device failures rapidly. When a Pivot3 node or disk fails, all remaining nodes and disks contribute capacity to return the cluster to a fully-protected status very quickly. Also, since Pivot3 combines all the disk, cache, and networking resources into a single aggregated storage pool, the overall performance impact is reduced should one (or more) of those devices or drive failures. Allied with this approach, the use of erasure coding increases the protection level of a volume to be able to tolerate an entire node failure and an additional drive failure, yet still maintain production availability and performance.

In contrast, vSAN availability is based in part on physical array constructs such as disk-based RAID groups. Thus, volumes get rebuilt onto the limited set of disks that make up a standard RAID set, not across all disks in the cluster. One outcome of vSAN’s disk-based RAID groups is that a VxRail cluster requires a minimum of five (5) nodes to keep a volume online in the face of just two simultaneous failures. (Figure 4.)

**FIGURE 4: Failures to Tolerate & Nodes Required Summary Comparison**

Failures to Tolerate	Number of Pivot3 Nodes Required	Number of VxRail Nodes Required*
1	3	3 (RAID-1) 4 (RAID-5)
2	3	5 (RAID-5) 6 (RAID-6)
3	3	7 (RAID-1)

\*VMware Validated Design on VxRail requires a minimum of 8 VxRail nodes including at least 4 nodes in the VM workload cluster plus a dedicated 4-node vSphere management cluster.

Similarly, vSAN has a structure called a disk group that consists of one cache drive plus one to seven capacity drives. If the cache drive fails—or any member of a vSAN disk group that has compression and deduplication enabled fails—the entire disk group fails and must be re-created elsewhere in the cluster. This extends the time that data may not be fully protected and makes vSAN capacity planning much more complex.

VMware’s [vSAN capacity planning guidelines](#)<sup>1</sup> include the following cautions. In addition to sacrificing 25% to 75% of raw capacity for RAID and replication, VMware advises vSAN users to...

1. “Keep at least 30 percent unused space to prevent vSAN from rebalancing... The rebalance operation might impact the performance of applications.”
2. “Plan extra capacity to handle any potential failure or replacement of capacity devices, disk groups and hosts... When a flash cache device fails or is removed, vSAN recovers the components from the entire disk group.” *(In plain English: The cache SSD is a single point of failure in every disk group. So, in addition to the 30% reserve, also reserve capacity that matches the largest disk group in the cluster.)*
3. “Reserve extra capacity to make sure that vSAN recovers components after a host failure or when a host enters maintenance

mode.” *(Reserve capacity equal to the largest node to ensure recovery without data loss.)*

4. “Provide enough temporary storage space for changes in the vSAN VM storage policy.” *(It is unclear what occurs if this condition is not met.)*
5. “If you plan to use advanced features, such as software checksum or deduplication and compression, reserve... approximately 6.2% capacity per device.” *(If you want to be able to verify that the data sent to vSAN was stored correctly and can be recovered if it wasn’t—that peace of mind will cost you another 6.2% of raw capacity.)*

Underestimating VxRail capacity reserves could negatively impact performance or result in data loss. If you decide to go with vSAN, you should most definitely overprovision storage! (Figure 5.)

**FIGURE 5: Availability & Reliability Summary Comparison**

	Pivot3	VxRail
Availability technology	Patented erasure coding	RAID plus 2-way, 3-way or 4-way Replication RAID 5/6 Erasure Coding
Usable vs raw capacity efficiency ratio	Up to 82%	25% to 50%
Premapped data blocks and metadata maps, data for each volume dispersed across all disks		

### KEY QUESTIONS TO ASK:

- What percentage of raw capacity will be wasted to meet my availability requirements? How much extra cost will this wasted capacity impose across the lifecycle of this infrastructure?
- What happens to performance and availability when a cluster node or disk fails?
- How many simultaneous failures can the cluster absorb and still serve up data?
- How many simultaneous failures can the cluster absorb without losing any data?
- How quickly can the solution self-heal to a fully protected state during node or disk failures?

## Pivot3 Security Extends into the Architecture

Pivot3 and VxRail are compliant with international Common Criteria<sup>2</sup> and FIPS 140-2<sup>3</sup> security standards. Both support data at rest encryption and the KMIP<sup>4</sup> security key management standard. However, Pivot3’s security is integrated into its Intelligence Engine’s policy-based management paradigm to simplify security management even at scale. (Figure 6.)

**FIGURE 6: Security Summary Comparison**

	Pivot3	VxRail
<b>Common Criteria Certification</b>	EAL 2+	EAL 2+
<b>Data at Rest Encryption (DARE)</b>	Software	Software
<b>FIPS 140-2</b>	✓	✓
<b>KMIP</b>	✓	✓

**KEY QUESTIONS TO ASK:**

- Does the solution comply with relevant security standards?
- How does the solution integrate with common enterprise security management infrastructure to enable management at scale?

**Pivot3 Flexible, Non-disruptive Scaling Yields Agility and Low TCO**

Early generations of HCI provided granular scalability, but not much flexibility in that all appliances in a cluster had to be identical. Thus, CPU and storage were added in lockstep whenever either resource became constrained. Early adopters accepted this inherent overprovisioning in order to gain the benefits of HCI. This is no longer the case.

While most HCI solutions now provide some flexibility in how they scale, meaningful differences remain. Pivot3 HCI clusters flexibly expand to meet virtually any data center requirement. In addition to the standard node, Pivot3 offers compute nodes, GPU-enabled nodes, all-flash storage nodes and hybrid storage nodes. The beneficial implications of this flexibility include:

- Customers can configure each new node to precisely address the customer’s needs without overprovisioning—and paying for—unnecessary performance resources.

**Notes**

1. "Planning Capacity in vSAN - VMware Docs." 6 Sep. 2018, <https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.vsan-planning.doc/GUID-581D2D5C-A88F-4318-A8B3-5A5F343F1247.html>. Accessed 25 Jan. 2019.
2. Common Criteria – Common Criteria for Information Technology Security Evaluation is an international standard (ISO/IEC 15408) for computer security certification. Common Criteria provides assurance that the process of specification, implementation and evaluation of a computer security product has been conducted in a rigorous and standard and repeatable manner at a level that is commensurate with the target environment for use. (<https://www.commoncriteriaportal.org/products/>)
3. FIPS – The FIPS 140-2 standard is an information technology security approval program for cryptographic modules produced by private sector vendors who seek to have their products certified for use in government departments and regulated industries (such as financial and health-care institutions) that collect, store, transfer, share and disseminate sensitive but unclassified (SBU) information.
4. Key Management Interoperability Protocol (KMIP) is protocol for communication between encryption systems and a broad range of new and legacy enterprise applications, including email, databases, and storage devices. KMIP enhances data security while reducing expenditures on multiple products and enables centralized encryption key management.

**About DCIG**

DCIG empowers the IT industry with actionable analysis that equips individuals within organizations to conduct technology assessments. DCIG delivers informed, insightful, third party analysis and commentary on IT technology. DCIG independently develops and licenses access to DCIG Buyer’s Guides and the DCIG Competitive Intelligence Platform. It also develops sponsored content in the form of blog entries, executive white papers, podcasts, competitive intelligence reports, webinars, white papers, and videos. More information is available at [www.dcig.com](http://www.dcig.com).

**FIGURE 7: Flexible Scaling Summary Comparison**

	Pivot3	VxRail
<b>Variety of node types and form factors</b>	✓	✓
<b>Clusters of mixed node types</b>	✓	Not recommended. First 3 nodes must be identical.
<b>Multiple server vendors</b>	✓	●
<b>Non-disruptive node addition</b>	✓	✓
<b>Non-disruptive node removal</b>	✓	✓
<b>Software-only licensing</b>	✓	●
<b>HCI running on public cloud infrastructure</b>	Pivot3 Cloud Edition (AMI in AWS)	VMware Cloud on AWS

- Pivot3 is equally effective for remote office and branch office (ROBO) applications, business continuity and disaster recovery (BC/DR), high-end video surveillance and enterprise SAN. (Figure 7.)

**KEY QUESTIONS TO ASK:**

- Can nodes be added and removed without forcing a restart of the cluster?
- What constraints does the solution impose on cluster designs?
- How many nodes are required to provide a highly available infrastructure solution?

**DCIG’s Conclusion:**

Pivot3’s intelligent policy-based management, sophisticated QoS, advanced erasure coding and flexible non-disruptive scaling elegantly converge to deliver an enterprise-class hyperconverged infrastructure solution. Taken together, these capabilities provide a compelling case for placing the Pivot3 Datacenter Series at the core of next generation data centers. ■

