Not all hyperconvergence is created equal

Pivot3's architecture is fundamentally different than other HCI vendors. Through erasure coding, Pivot3 distributes data across the system, ensuring maximum performance and resiliency are always maintained without resorting to hardware-based replication. Pivot3's erasure coding-based software aggregates all available resources into a unified pool that can be accessed by all virtual machines from any node in the system, which inherently improves the efficiency and performance of all resources. When new nodes are added, the entire environment benefits from the additional performance and capacity. The result is better performance, more usable capacity, and higher availability than replication-based HCI vendors to keep your Total Cost of Ownership low while reducing complexity and risk.

Most other HCI vendors rely on replication-based storage architectures, which significantly impact usable storage capacity, availability and performance. When these vendors' systems create VMs, associated data is stored on a single node and relies on a file system to share storage across nodes so VMs can move and still access that data. This approach not only adds extra overhead, it also can create performance bottlenecks. Pivot3 employs traditional block file protocols to avoid impacting performance. When these systems add new nodes, data cache, IOPS and bandwidth are not aggregated into a unified pool for use by all VMs, they remain node-local, resulting in inefficiencies that extend to all high-demand applications.

Pivot3 with Erasure Coding

Rather than resorting to replication to protect data, Pivot3's erasure coding enables data distribution to multiple media types across nodes to maximize usable storage space and ensure more robust fault tolerance.

Data is written once across the entire environment, which protects users from multiple failure events, speeds recovery with low system impact, and keeps data efficiency very high. Block-based SAN storage leverages internodal communications without the need for a metadata server, removing performance bottlenecks without impacting availability

VIRTUAL SAN

All data is wide-striped across the system, aggregating all resources into one shared pool to increase performance for all virtual machines.

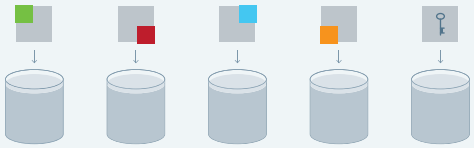

Utilize up to 94% of total raw storage capacity. Selective compression for applications can double system capacity.

vs.

Replication-Based HCI

Other HCI vendors run VMs locally on the resources of the underlying server, they only work with other nodes to make copies of data but do not aggregate resources across nodes.

Replication-based protection schemes consume 2–3X more storage capacity to write the same amount of data. Scale-out file systems require metadata server to manage data placement, which can be a performance bottleneck, as well as a single-point of failure.

VMs are not automatically load balanced across the system, causing performance and management issues because the default is to disable features like VMware DRS.

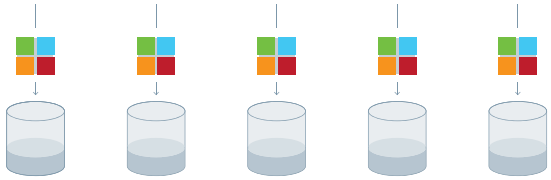

Usable capacity drops with each added level of redundancy. At one copy, only 50% of deployed storage is available for use, dropping below 33% with two.

vs.

Pivot3 with Erasure Coding

Erasure coding protects from a greater number of failures within the system.

Maximum failures before unavailability

5 Drives
or
2 Drives and 1 Node

SIX
9'S AVAILABILITY

**Using half as many nodes*

Even in the smallest configurations, there is less than a 15% degradation in performance during a failure scenario. HDD failure is not a factor — predictive sparing prevents failure — when a drive begins to fail, system begins to copy data non-disruptively. A high resilience (up to five drives) means you don't have to rebuild immediately. There is NO danger of permanent data loss.

<10%
OVERHEAD

vSTAC OS creates a lean operating environment, with an extremely low overhead to devote more compute power to business applications.

Replication-Based HCI

Losing more than three drives or two drives and a node results in critical data loss.

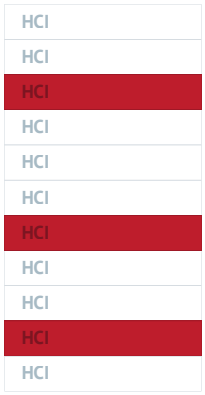

Maximum failures before critical data loss

2 Drives
or
1 Node

THREE
9'S AVAILABILITY

Rebuilds and recovery activities can negatively impact main functions and foreground applications. VM must migrate to new appliance and data must be copied over to the appliance where the VM now operates. HDD failures must rebuild. During that time, CPU performance and available capacity is relegated to rebuild. Because it is replication-based, you only get two copies. Drive loss means you are one drive away from permanent data loss during a rebuild.

40%
OVERHEAD

Emphasis on deduplication and compression add overhead, consuming precious CPU cycles. Consume 40% or more of system resources just to operate the appliance.

Pivot3 Scalar Erasure Coding 7.X vs. Erasure Coding 1.0

Almost all other HCI providers have long touted the superiority of a replication-based approach, but last June another HCI provider announced that it would be updating it's software to include a form of erasure coding to improve storage effectiveness. At version 7.X, Pivot3 has been developing, refining and perfecting its patented Scalar Erasure Coding for years, delivering optimal results for over 2,000 customers in 53 countries around the world and continues to outperform the competition in terms of performance, resiliency, and storage efficiency. Take a look at how the two stack up:

CAPABILITY	Pivot3	HCI COMPETITOR
Erasure Coding Version	**7** SINCE 2005	**1** SINCE 2015
Erasure Coding Always On (Inline)	**Yes** OPTIMIZED AT WRITE, PROTECTION ACROSS ALL NODES	**No** SOLUTION REQUIRES DATA REPLICATION OVERHEAD TO CPU AND STORAGE COPIES
Erasure Coding Batch Process	**No**	**Yes**
Erasure Coding Overhead (CPU)	**Less than 3%** ALWAYS ON	**Greater than 20%** ON TOP OF DISTRIBUTED FILE SYSTEM OVERHEAD WHICH TYPICALLY IS 45% AT 3 NODES AND UP TO 58% AT 10 NODES
Erasure Coding Applies to	**All Data** DELIVERS HIGH USABLE STORAGE AND LOW OVERHEAD	**Cold Data** ONLY DATA THAT HAS "COOLED" FOR 1 HOUR AND DONE AS A BATCH OFFLINE PROCESS
Does Solution Require Data Copies/ Replication to Protect Data	**No** DATA REPLICATION IS INEFFICIENT AND CONSUMES 3X STORAGE	**Yes** REDUCES USABLE SPACE BY UP TO 80% AND MAKES USE ON LIMITED STORAGE FORM FACTORS SUCH AS FLASH/ BLADES UNFEASIBLE DUE TO LIMITED CAPACITIES ON THOSE DEVICES
Does the solution require data to be on the node (local copy) for processing, or is data striped across all nodes	**No** AFTER A VMOTION ACTIVITY, DATA REMAINS IN PLACE AND DOES NOT MOVE TO FOLLOW THE VM BECAUSE DATA IS WIDE-STRIPED AND AVAILABLE ACROSS ALL NODES.	**Yes** LOCAL HYPERCONVERGENCE LIMITATION MEANS DATA NEEDS TO BE MOVED PROXIMATE TO THE COMPUTE NODE FOR PERFORMANCE EFFICIENCY, INCREASING STORAGE AND CPU OVERHEAD DUE TO REPLICATION
Percentage of Data Erasure Coding Applies to	**100%**	**Less than 10%** ONLY ON "COOLED" DATA AS AN OFFLINE PROCESS, RE-DUCING ACTUAL % OF ERASURE-ENCODED DATA; CAN NOT REPLACE REPLICATION AND ERASURE CODING CANNOT BE PRIMARY STORAGE METHOD
Effective Usable Storage (Percentage) at 12 Nodes, Protection Level 1	**94%**	**46%**
Effective Usable Storage (Percentage) at 16 Nodes, Protection Level 3	**81%**	**33%**
Effective Usable Storage (Percentage) at 12 Nodes, Protection Level 5	**76%**	**16%**
Projected Usable Storage Across Erasure Coding or with Competitor's Replication + Erasure Coding at Protection Level 3	**94%**	**40%**

Appendix of Comparison Examples

PROTECTION LEVELS

PIVOT3 ERASURE CODING

PROTECTION LEVEL	EFFICIENCY*
Scalar Erasure Coding Lvl.5 (12 Nodes) 5 simultaneous disk events or 2 drives and 1 node	76%
Scalar Erasure Coding Lvl.3 (12 Nodes) 3 simultaneous disk events or 1 drive and 1 node	83%
Scalar Erasure Coding Lvl.1 (12 Nodes) 1 disk event or 1 drive and 1 node	94%

** does not include sparing*

REPLICATION-BASED HCI (best match)

PROTECTION LEVEL	EFFICIENCY
3 Replicated Copies 3 simultaneous disk events or 1 node	25%
2 Replicated Copies 2 simultaneous disk events or 1 node	33%
1 Replicated Copy 1 disk event or 1 node	50%

STORAGE EFFICIENCY

	EXAMPLE NO. 1		EXAMPLE NO. 2	
	PIVOT3	REPLICATION-BASED HCI	PIVOT3	REPLICATION-BASED HCI
Usable TB Target	100 TB	100 TB	200 TB	200 TB
Raw Capacity Per Node	24 TB	24 TB	24 TB	24 TB
Required Nodes	6	9	11	26
Incremental Nodes %	-	50%	-	136%
Raw Capacity Purchased	144 TB	216 TB	240 TB	624 TB
Storage Efficiency	81%	50%	80%	33%
Storage Protection	Lvl.1	1 Replicated Copy	Lvl.3	2 Replicated Copies
Usable Storage	116.3 TB	108 TB	211.8 TB	207.8 TB

About Pivot3

Pivot3 delivers Dynamic Hyperconvergence through innovative technologies to deploy intelligent compute and storage infrastructures. Pivot3's solutions maximize resource utilization, deliver sustained business services and guarantee application performance levels to improve data center simplicity and drive down complexity and cost. Today, Pivot3 has more than 2,000 customers in 53 countries and has successfully deployed over 16,000 mission- and business-critical infrastructures in multiple industries such as video surveillance, health-care, government, transportation, education, gaming, entertainment and retail.

Learn more about Pivot3's unique innovation and technological differentiation.
Visit Pivot3.com/Why-Pivot3/