


Hyperconverged Infrastructure Won the Battle.

Then it created the next problem.



HCI removed the storage handshake. The handshake. The win was real.


01 — THE ROUND TRIP

Storage I/O stopped crossing a separate fabric to a separate array by default.
Storage processing moved into software across the nodes.

02 — THE TICKET LOOP

Request, carve, mask, zone, rescan, format. Every step a handoff between between teams. HCI replaced the handoff with a policy.

The node became the unit of architecture. architecture.



Six concerns. One object. That is why HCI was simple, and why it won.

The uniform node breaks on four axes, not axes, not one.

GPU

REFRESH LIFECYCLE

Accelerators turn over on a different clock than storage.

STATEFUL K8S

STORAGE AUTHORITY

Persistent volumes can bring a second second storage control plane.

EDGE

FAILURE TOLERANCE

Availability survives. N+1 capacity often often does not.

CAPACITY-DENSE

ECONOMICS

Storage grows. Compute gets bought bought anyway.

Economics shows up on a quote. The other three show up later.

THE BOUNDARY IS RETURNING

The market is not abandoning convergence. convergence.

It is selectively restoring separation where uniform nodes no longer fit the workload.

CONTROL PLANE — STILL CONVERGED

COMPUTE NODES

STORAGE TIER

DATA PLANE — SEPARATING AGAIN

NUTANIX

FY2026 10-K: qualified external storage integrations


VMWARE

vSAN storage clusters: disaggregated deployment option

VCF 9.0

Separate VMkernel traffic for storage-cluster clients

The handshake comes back. The other side of it doesn't.



The boundary is returning faster than the expertise that used to govern it.

Name the seam owner before you separate. separate.

WORKLOAD SIGNAL	POSTURE	SEAM OWNER
Uniform VM estate	Stay converged	No seam — node owner stays accountable
Storage outgrows compute	Storage cluster or external array	Named owner: fabric, firm ware, capacity, escalation
GPU refresh diverges	Separate accelerator nodes	Named owner for the data path
Stateful K8s at scale	Dedicated PV storage	Owner per operation: snapshots, replication, restores
Edge footprint	Stay converged	Central platform team ; site failure model documented

When storage and virtualization releases fall out of compatibility, who owns the decision to decision to proceed?

Hyperconverged infrastructure collapsed the storage boundary.

It never eliminated responsibility for what lived on either side of it.

Full post: rack2cloud.com/hyperconverged-infrastructure-next-problem

Follow The Architect for more.