



>\_ AHV PRE-MIGRATION CHECKLIST

## 9 Validation Steps Before Your First Production Cutover

Part of the Post-Broadcom Migration Series | rack2cloud.com

>\_ READINESS - SIZING & ARCHITECTURE

### 1 CVM Sizing Against Your Actual Workload Profile

Not the minimum, not the default. The HCI Migration Advisor runs automated readiness checks — detecting ISOs, snapshots, and hardware compatibility — and maps your I/O profile against CVM sizing before you commit to node count.

### 2 CVM NUMA Alignment

Large CVM vCPU allocations can cross NUMA boundaries on multi-socket hosts, introducing remote memory latency. Validate NUMA topology, confirm CVM vCPU fits within a single NUMA domain, and verify memory locality. Critical on AMD EPYC hosts.

### 3 Storage Network Validation

Every write is replicated across nodes before acknowledgement. Validate: 25GbE minimum for NVMe nodes (10GbE is insufficient), jumbo frames enabled end-to-end, ToR switch buffer depth adequate for write burst profile. A network bottleneck looks exactly like CVM CPU saturation.

### 4 RF Policy Impact on Stargate Load

RF2 doubles every write. RF3 triples it. If your workload is write-heavy and you default to RF3 for all VMs, your Stargate CPU baseline is 50% higher than an equivalent RF2 deployment. Model RF policy against your write I/O profile before migration.

### 5 Snapshot Policy Review

Deep snapshot trees drive sustained Curator and Cassandra overhead on the CVM. Audit existing snapshot depth, backup schedule frequency, and retention policies before migration. A 30-day snapshot chain moves with the VM — and so does the Curator load it generates.

>\_ RF WRITE AMPLIFICATION REFERENCE

RF Policy	Write Amplification	Stargate CPU Impact
RF2	2×	Baseline
RF3	3×	+50% above RF2 baseline



## >\_ READINESS — CAPACITY & OPERATIONS

### Cluster Rebuild Time Estimation

If a node fails, surviving CVMs must rebuild the failed node's data share — a sustained I/O and CPU spike running parallel to VM workloads. Estimate before migration:

$$6 \quad \text{cluster\_rebuild\_time} \approx \text{dataset\_size} / \text{replication\_network\_bandwidth}$$

A 20TB dataset on 25GbE at 50% utilisation = multi-hour elevated load event. Know the number.

### Effective VM Pool After CVM and Hypervisor Reservation

Recalculate VM density using effective capacity, not raw host specs. Subtract CVM reservation (typically 12 vCPU / 32GB+) and hypervisor overhead (~2 cores / 8GB) from raw host resources before any VM placement math.

7

### N+1 CVM Headroom

Model your cluster with one node's CVM offline. If surviving CVMs exceed 50% CPU utilisation at peak under that scenario, add a node before migration. The 50% threshold is the buffer that absorbs N+1 load without degrading guest performance.

8

### Monitoring Baseline Established Before Cutover

Configure Prism Central alerts for CVM CPU threshold, remote I/O percentage, and CVM memory ballooning before the first production VM moves. Establish the baseline during low-stakes migration waves so anomalies stand out clearly during high-stakes cutovers.

9

## >\_ CONTROLLER OVERHEAD — PLATFORM COMPARISON

Platform	Controller Architecture	Overhead Model	Predictability
Nutanix AHV	CVM per node	Fixed CPU + memory reservation	High — explicit upfront
VMware vSAN	Kernel module (vmkernel)	Minimal CPU, kernel scheduler	Medium — hides under load
Ceph / Proxmox	OSD daemon per disk	CPU scales with drive count	Medium — scales with density

### >\_ KEY INSIGHT

Nutanix pays the controller tax up front — explicitly, in the capacity model. Ceph distributes it across disks. vSAN hides it in the kernel scheduler.

### >\_ TOOL: HCI MIGRATION ADVISOR

#### Model your controller tax before you commit to node count.

Automated readiness checks, CVM sizing, N+1 headroom model.  
[rack2cloud.com/vmware-to-hci-migration-advisor/](https://rack2cloud.com/vmware-to-hci-migration-advisor/)