SOLUTION BRIEF

SUPERMICRO, PORTWORX AND RED HAT BUILD A RESILIENT AND COST-EFFECTIVE EDGE SOLUTION

A Co-Engineered Solution with Supermicro Systems, Red Hat OpenShift, and Portworx by Everpure

Supermicro SYS-121C-TN10R and AS -E300-14GR Servers

TABLE OF CONTENTS

Executive Summary: The Rise of the Intelligent Edge

The demand to extend applications, AI inferencing, and data processing to the edge has never been greater. From retail stores and industrial sites to 5G towers and healthcare clinics, organizations are creating and consuming massive volumes of data closer to where they originate — enabling real-time insights, ultra-low latency, and operational efficiency that centralized cloud architectures cannot deliver.

Edge deployments impose severe constraints: tight physical space, restricted power and cooling, and the need to remotely manage hundreds or thousands of sites without local IT staff. Yet high availability (HA) remains non-negotiable. Mission-critical systems — point-of-sale platforms, factory robotics, 5G network processing — must survive hardware failure without disruption or data loss.

Traditionally, achieving full HA resilience required at least three identical server nodes to establish a control plane and a storage quorum. This solution brief describes a production-validated, co-engineered reference architecture that combines Supermicro purpose-built edge servers, Red Hat OpenShift Platform Plus, and Portworx by Everpure container-native data services — delivering the Two-Node with Arbiter (TNA) topology for full three-node HA resilience at a significantly lower cost and footprint.

Core Solution Efficiency Gains

- 33% Hardware Reduction: Eliminates a full-sized server from the per-site bill of materials, without sacrificing HA resilience.
- Up to 65% Savings: Moving to the Enterprise standard for container management, Red Hat OpenShift Platform Plus, and reducing active node counts bypasses legacy per-core hypervisor taxes.
- 30% Lower Power Draw: Minimized server footprint combined with ultra-efficient Supermicro Titanium PSUs reduces operational energy overhead.
- Data High Availability: Continuous synchronous block-level replication guarantees no data loss during any abrupt node failure.

Solution Architecture: Asymmetric Quorum Design

The TNA architecture operates logically as a three-node cluster for quorum purposes, but uses an asymmetric hardware configuration to dramatically reduce costs. Two full-performance Supermicro data nodes run all workloads and Portworx by Everpure, replicated storage. A single lightweight arbiter node provides quorum without running any workloads or storing data volumes.

Two Data Nodes — Supermicro Performance Servers: Full-performance 1U or 2U servers hosting the Red Hat OpenShift control plane, all containerized and VM workloads (via OpenShift Virtualization), and Portworx synchronously replicated NVMe storage volumes.

One Arbiter Node — Supermicro SYS-111C-NR or AS -E300-14GR: A single lightweight compact server running only the third OpenShift etcd instance and the Portworx KVDB metadata engine. The arbiter stores no data volumes, runs no user applications, and requires minimal compute resources — and zero heavy software licensing.

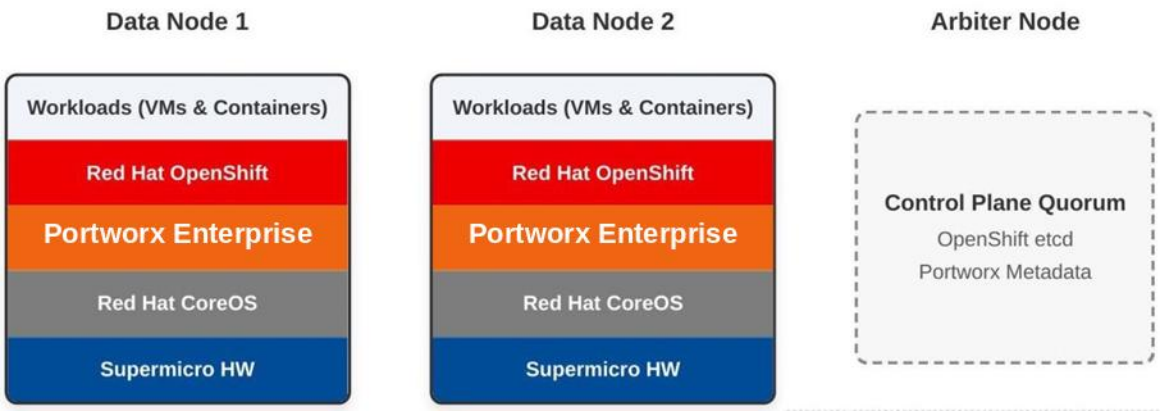

Figure 1 - TNA Architecture — two Supermicro data nodes host all workloads and synchronously replicate storage; a compact arbiter maintains quorum for OpenShift etcd and Portworx KVDB without storing any data.

Network Quorum, Latency, and Split-Brain Protection

To maintain data integrity at the edge, the storage and control networks must handle transient network drops safely. The TNA topology provides three key safeguards:

- Scheduling Restrictions: Red Hat OpenShift Platform Plus natively prevents application container and VM workloads from executing on the Arbiter node, keeping its processing dedicated to control-plane health.
- Synchronous Replication Requirements: Portworx writes data simultaneously to both data nodes before confirming I/O success.
- Split-Brain Mitigation: If the link between the two data nodes breaks but both can still reach the Arbiter, a majority is preserved (2 of 3 votes). The isolated data node is automatically excluded from storage operations.

Supermicro Hardware Infrastructure

Supermicro delivers data center-class computing power in compact, ruggedized form factors tailored for space- and power-constrained edge conditions. To accommodate varied edge deployment profiles, Supermicro offers three validated sizes for data nodes — Small, Medium, and Large — while the lightweight arbiter node remains consistent across all tiers.

Three data node tiers cover the full range of edge deployment profiles — from compact far-edge closets to dense AI-ready industrial racks — while the arbiter remains lightweight across all configurations.

Supermicro's hardware advantages are purpose-built for always-on edge deployments. Titanium-rated (96%+ efficiency) redundant power supplies minimize energy draw and heat generation at every site — critical when running 24/7 in space- and power-constrained locations. All data node tiers feature NVMe-native storage architectures with PCIe 5.0 connectivity, delivering the low-latency I/O throughput required by Portworx synchronous replication without an external SAN. IPMI 2.0 BMC and SuperDoctor 5 are standard across every platform, enabling full remote monitoring, BIOS/firmware lifecycle management, and predictive failure alerting from a central console — eliminating the need for on-site technicians for routine maintenance. Supermicro servers are also MIL-grade vibration- and shock-tested for industrial and far-edge environments and support tool-less Field Replaceable Units (FRUs) to minimize mean time to repair (MTTR) when a component fails. Together, these platform advantages make Supermicro the ideal hardware foundation for a production-grade, unattended TNA deployment at scale.

Specification	SMALL Entry / Far Edge **AS -1114S-WN10RT**	MEDIUM Retail / 5G / Mid-range **SYS-121C-TN10R**	LARGE Dense AI / Industrial **SYS-221U-TN24R**	ARBITER All tiers **SYS-111C-NR / AS -E300-14GR**
Form Factor	**1U Rackmount**	**1U Rackmount**	**2U Rackmount**	1U / Mini-1U (<9")
CPU	Single AMD EPYC 8–32 cores	Dual Intel Xeon 16–52 cores	Dual Intel Xeon 32–104 cores	Intel Xeon / AMD EPYC 4005 4–16 cores
Memory	32–256 GB DDR5	64–512 GB DDR5	128 GB–2 TB DDR5	8–32 GB DDR5
NVMe Storage	8–10 hot-swap bays	10 hot-swap bays	Up to 24 hot-swap bays	M.2 NVMe only (KVDB)
GPU/AI Ready	—	NVIDIA L4/T4	NVIDIA L4/T4	—
Network	Dual 25 GbE	Dual 25 GbE	Dual 25 GbE	1–10 GbE
Product Links	supermicro.com →	supermicro.com →	supermicro.com →	supermicro.com →

Architectural Note: Heterogeneous, mixed-processor clusters are fully supported in this TNA validation. AMD-powered data nodes (Small Tier) can coexist with an Intel-powered SYS-111C-NR arbiter. The arbiter processes control-plane quorum metadata only and runs no user workloads, avoiding any cross-architecture binary compilation issues.

Arbiter Node Options

Two validated arbiter platforms serve all three data node tiers — selected based on available rack space and environmental conditions:

SYS-111C-NR — Primary Arbiter

1U Rack · Intel Xeon · Red Hat OCP Certified

- Standard 1U depth — fits any server rack
- Intel Xeon — Red Hat OpenShift certified
- NVMe / SSD for KVDB and etcd metadata only
- IPMI 2.0 BMC + SuperDoctor 5 remote management

AS -E300-14GR — Compact Arbiter

Mini-1U · AMD EPYC 4005 · 8.89" depth · 65W TDP

- Ultra-compact: 8.89" (225 mm) depth — fits any closet
- AMD EPYC 4005 up to 16 cores · 65W TDP
- Up to 192 GB DDR5 ECC · M.2 NVMe + 2.5" SATA
- Wide-temp + fanless option for harsh environments

Production-Validated Software Stack

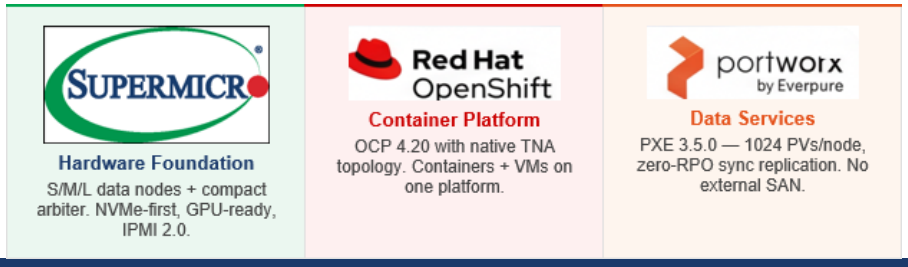

Software Component	Validated Version
Supermicro BIOS — Data Nodes (all tiers)	Latest validated per platform
Supermicro BIOS — SYS-111C-NR / AS -E300-14GR (Arbiter)	Latest validated (Xeon / EPYC 4005)
Supermicro BMC / IPMI Firmware + SuperDoctor 5	IPMI 2.0 — latest stable release
Portworx Enterprise	PXE 3.5.2
Red Hat Enterprise Linux CoreOS	RHCOS 9.6
Red Hat OpenShift Platform Plus	OCP 4.20 (GA) — native TNA topology

Table 1 - Production-validated software stack. Green = Supermicro hardware lifecycle components.

Operational Walkthrough: Failover, Recovery, and Day-2

High-Availability Failover and RTO Sequence

When a primary data node experiences an abrupt hardware failure — localized power loss, motherboard fault — the cluster manages recovery automatically in four stages:

- Failure Detection: Data Node 1 drops offline. Heartbeat checks fail across both the OpenShift etcd and Portworx KVDB control paths.
- Quorum Verification: Data Node 2 and the Arbiter maintain a functional majority (2 of 3 etcd/KVDB votes), keeping the cluster control plane write-active.
- Storage Transition: Portworx automatically flags storage volumes on Data Node 2 as the primary active writable path. Because of continuous synchronous mirroring, there is zero data loss — RPO = 0.
- Workload Recovery: OpenShift detects the lost node capacity and immediately restarts container pods and Virtual Machines (via Red Hat OpenShift Virtualization) on Data Node 2 — restoring application operations in seconds. RTO is measured in seconds, not minutes.

Day-2 Operations and Rolling Updates

Executing software upgrades or security patches across hundreds of remote edge sites is complex. The TNA architecture natively supports zero-downtime rolling updates across all three data node tiers:

- Evacuation: Red Hat Advanced Cluster Management (ACM) triggers a maintenance window for Data Node 1. Active workloads and running VMs are live-migrated to Data Node 2 using OpenShift Virtualization.

- Upgrade and Reboot: Data Node 1 is upgraded — host patches, Supermicro BIOS/BMC firmware updates via SuperDoctor 5, or Portworx storage updates — and rebooted. The Arbiter and Data Node 2 maintain an active control plane quorum and host all production traffic.
- Re-Integration: Data Node 1 boots back into the cluster. Portworx triggers a fast catch-up block sync to update mirrored NVMe drives. The process then repeats for Data Node 2 — delivering a fully patched cluster with zero application downtime.

Portworx StorageCluster Configuration

The following StorageCluster CRD establishes the TNA topology within Portworx Enterprise, per Portworx and Red Hat's validated reference spec. The arbiter node's storage block omits a devices list entirely (only a kvdbDevice and useAll: false) — ensuring it participates solely as a metadata quorum witness and never absorbs application data storage paths.

Production Best Practice: Use persistent device symlinks (/dev/disk/by-id/*) rather than /dev/nvmeXnY paths. The kernel device enumeration order can change across reboots on Supermicro platforms; persistent IDs ensure that Portworx always attaches to the correct physical NVMe drive.

```
kind: StorageCluster
apiVersion: core.libopenstorage.org/v1
metadata:
  name: px-cluster-smci-tna
  namespace: portworx
  annotations:
    portworx.io/is-openshift: "true"
    portworx.io/run-on-master: "true"
    portworx.io/misc-args: "-T px-storev2 -rt_opts small_conf=1"
    portworx.io/preflight-check: "skip"
    portworx.io/health-check: "skip"
spec:
  image: portworx/oci-monitor:3.5.2
  kvdb:
    internal: true
  nodes:
    # Data Node 1 — Supermicro Performance Platform (Small / Medium / Large tier)
    - clusterDomain: smci-worker-1
      selector:
```

```
      nodeName: smci-worker-1
    storage:
      devices:
        - /dev/disk/by-id/nvme-SMCI_DATA_DRIVE_DATA01
      systemMetadataDevice: /dev/disk/by-id/nvme-SMCI_STABLE_META01
  # Data Node 2 — Supermicro Performance Platform (same tier as Node 1)
  - clusterDomain: smci-worker-2
    selector:
      nodeName: smci-worker-2
    storage:
      devices:
        - /dev/disk/by-id/nvme-SMCI_DATA_DRIVE_DATA02
      systemMetadataDevice: /dev/disk/by-id/nvme-SMCI_STABLE_META02
  # Arbiter Node — Supermicro SYS-111C-NR or AS -E300-14GR
  - clusterDomain: witness
    selector:
      nodeName: smci-arbiter
    storage:
      kvdbDevice: /dev/disk/by-id/nvme-SMCI_M2_METADATA_ONLY
      useAll: false        # NO application data volumes on arbiter
placement:
  tolerations:
    - {effect: NoSchedule, key: node-role.kubernetes.io/arbiter}
    - {effect: NoSchedule, key: node-role.kubernetes.io/master}
    - {effect: NoSchedule, key: node-role.kubernetes.io/control-plane}
csi:       {enabled: true}
stork:     {enabled: true}
autopilot: {enabled: true}
monitoring:
  telemetry:  {enabled: true}
  prometheus: {enabled: true, exportMetrics: true}
```

Key Solution Benefits

High Availability

Portworx synchronously replicates data at the block level between the NVMe drives of both data nodes. Upon a node failure, OpenShift automatically restarts all workloads on the surviving node, which already has a mirrored copy. RPO = zero, RTO = seconds, regardless of data node tier.

Consolidated Compute and Storage

Red Hat OpenShift offers a unified platform to run traditional VMs and modern, containerized workloads side by side. By allowing VMs to share the same high-efficiency Supermicro hardware, you eliminate the cost and complexity of a separate hypervisor infrastructure. Two data nodes cut per-site hardware, power, and cooling by 33% vs. a traditional 3-node cluster.

Up to 65% Licensing Savings

Consolidating open-source Kubernetes and reducing heavy node count from three to two delivers up to 65% savings in software licensing costs — plus 30% lower power draw via Supermicro Titanium PSUs — across every edge site in the deployment.

Central Management at Scale

Red Hat Advanced Cluster Management combined with Supermicro IPMI 2.0 BMC and SuperDoctor 5 provides a single pane of glass for deploying OS images, pushing firmware updates, and managing Kubernetes clusters across thousands of edge sites remotely.

Enterprise Use Cases

Manufacturing & Industrial

Vision quality inspection, predictive maintenance, and legacy AOI alongside containerized AI — on a single Supermicro Large-tier node.

AI / ML at the Edge

GPU inferencing on Large-tier SYS-221U or Medium-tier SYS-121C with NVIDIA L4/T4 via PCIe 5.0. Sub-ms latency, zero cloud round-trips.

Healthcare

Vital-signs monitoring, surgical imaging, and zero-RPO Portworx replication for mission-critical patient data on Red Hat OpenShift.

5G and Telecom Edge

Open RAN and 5G CNFs on Medium-tier Supermicro platforms validated for vRAN and o-RAN low-latency, high-throughput requirements.

Retail and POS Edge

Small-tier AS -1114S-WN10RT handles all POS, inventory, and AI-driven analytics. Compact AS -E300-14GR arbiter fits in the store.

Drones and Robotics

On-board inferencing for GPS-denied operations. Portworx manages sensor datasets; OpenShift orchestrates fleet management.

Conclusion

The Supermicro TNA-validated portfolio gives organizations the flexibility to deploy right-sized server infrastructure at any edge environment — from a compact Small-tier AS -1114S-WN10RT for space-constrained retail closets, to a dense Large-tier SYS-221U-TN24R with up to 24 NVMe bays for intensive industrial AI processing.

By pairing two performance data nodes with a highly efficient SYS-111C-NR or AS -E300-14GR arbiter, every deployment achieves robust three-node high-availability protection with a reduced two-node hardware bill of materials. Red Hat OpenShift delivers the enterprise Kubernetes foundation for running containers and legacy VMs side by side, eliminating hypervisor overhead. Portworx Enterprise provides synchronous block replication to maintain absolute zero RPO with no external SAN requirements.

The result: 33% less hardware per site, up to 65% lower software licensing costs, and a 30% reduction in power draw — making large-scale edge automation both technically and financially viable at any deployment scale.

Solution at a Glance

- Architecture: Three-node resilience, two-node hardware footprint — full HA without a third full-sized server.
- Small — AS -1114S-WN10RT: 1U, Single AMD EPYC,8–10 NVMe bays, Entry / Far Edge deployments.
- Medium — SYS-121C-TN10R: 1U, Dual Intel Xeon, 10 NVMe, GPU L4/T4, Retail / 5G / Mid-density.
- Large — SYS-221U-TN24R: 2U · Dual Intel Xeon · Up to 24 NVMe · GPU L4/T4 · Dense AI / Industrial.
- Arbiter — SYS-111C-NR / AS -E300-14GR: 1U or Mini-1U · M.2 NVMe only, No workloads, No data drives, IPMI managed.
- Platform: Red Hat OCP 4.20 on RHCOS 9.6 — containers + VMs on one platform, no separate hypervisor.
- Storage: Portworx PXE 3.5.2 — 1024 PVs/node, zero-RPO sync replication, live migration, no SAN.

For More Information

- Supermicro: www.supermicro.com
- Portworx: www.portworx.com
- Red Hat: https://www.redhat.com/en/hybrid-cloud-solutions

SUPERMICRO

As a global leader in high performance, high efficiency server technology and innovation, we develop and provide end-to-end green computing solutions to the data center, cloud computing, enterprise IT, big data, HPC, and embedded markets. Our Building Block Solutions® approach allows us to provide a broad range of SKUs, and enables us to build and deliver application-optimized solutions based upon your requirements. Visit www.supermicro.com

PORTWORX BY EVERPURE

Everpure (NYSE: P), formerly Pure Storage, is an American technology company headquartered in Santa Clara, CA, delivering an integrated storage and data management platform for the AI era. Its Portworx platform provides enterprise Kubernetes-native data services — including container-native storage, disaster recovery, backup, and data security — enabling mission-critical applications across on-premises, hybrid cloud, and edge environments. Visit www.everpuredata.com

RED HAT

Red Hat is the open hybrid cloud technology leader, delivering a trusted, consistent and comprehensive foundation for transformative IT innovation and AI applications. Its portfolio of cloud, developer, AI, Linux, automation and application platform technologies enables any application, anywhere—from the datacenter to the edge. As the world's leading provider of enterprise open source software solutions, Red Hat invests in open ecosystems and communities to solve tomorrow's IT challenges. Collaborating with partners and customers, Red Hat helps them build, connect, automate, secure and manage their IT environments, supported by consulting services and award-winning training and certification offerings. Visit www.redhat.com