



# Supermicro Systems Solution for Red Hat Hyperconverged Infrastructure for Virtualization (RHHI-V)

*New Compact, Simplified, Cost-Effective Compute/Storage Systems to Extend Red Hat-based IT Infrastructure*

## Executive Summary

Supermicro systems combined with Red Hat® Hyperconverged Infrastructure for Virtualization (RHHI-V) offers a simple, low-cost, and proven Hyperconverged Infrastructure (HCI) building blocks to run applications and data storage for edge and branch offices. These systems are designed to easily extend an existing data center with a Red Hat environment into the edge. Windows and Linux Server applications and virtual desktops are supported.

Benefits of Supermicro systems running RHHI-V include:

- o Time saving due to its simplicity and a modular scaling
- o Lower cost as all the HCI features and support are included in one price
- o Built from Supermicro's 27-years of proven industry system design coupled with Red Hat's comprehensive HCI software stack

Organizations in industries like energy, retail, banking, telecommunications, and the public sector increasingly rely on critical applications that must be deployed with limited space, budgetary constraints, and a scarcity of skilled IT staff. They must extend the reach of these applications from the data center core to the edge of the network. Supermicro system solutions running RHHI-V are designed for these environments.

In today's ever-changing digital landscape, organizations are searching to manage and deploy applications faster while reducing the total cost of ownership. As your IT grows, traditional infrastructures require added hardware, software, and resources with application growth. For organizations to remain relevant, IT teams need to remain efficient, effective and, responsive. HCI technology allows for simplicity and ease of deployment while lowering the total cost of ownership. With fewer systems to manage, hyperconverged infrastructures reduce deployment time, design time, and integration complexity. Supermicro and Red Hat have collaborated to create HCI for virtualization



Red Hat Hyperconverged Infrastructure for Virtualization



solutions designed to augment your IT infrastructure by merging compute, storage and, networking into a single easy-to-manage solution.

Red Hat Hyperconverged Infrastructure Virtualization and Supermicro Green IT servers provide a fully integrated HCI solution that can reduce deployment times from days to hours. This removes the complexity of having to manage traditional IT environments that encompass servers, storage, and networking switches while providing the automation of hardware and software to quickly deploy a fully operational system. Red Hat Hyperconverged Infrastructure Virtualization provides an open-source platform for deploying virtualized applications with the power of Supermicro’s industry-leading high-performance, high-efficiency, green computing server technology that provides advanced server solutions for the datacenter.

## TABLE OF CONTENTS

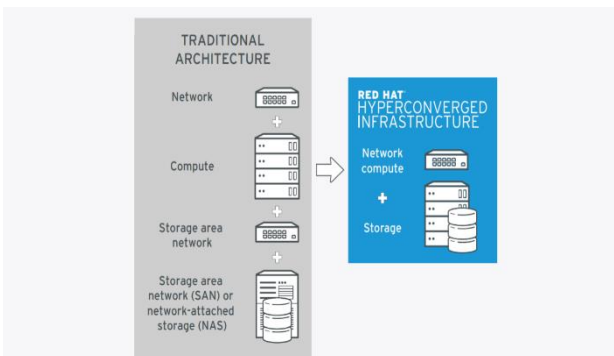
Introduction to Red Hat HCI Infrastructure for Virtualization Features .....	2
Value of Supermicro Solutions for Red Hat HCI Infrascture for Virtualization.....	3
Red Hat HCI Infrascture Architect Overview.....	4
Supermicro Solutions and Reference Architectures for Red Hat RHHI-V .....	5
Introduction to Supermicro BigTwin: Industry’s Highest Performing Twin Multi-Node System .....	6
The 5 <sup>th</sup> Generation Twin Architecture .....	7
Suggested Applications .....	8
Summary .....	9

## SUPERMICRO

Supermicro is a global leader in high performance, green computing server technology and innovation. We provide our global customers with application-optimized servers and workstations customized with blade, storage, and GPU solutions. Our products offer proven reliability, superior design, and one of the industry’s broadest array of product configurations, to fit all computational need.

# Introduction to Red Hat Hyperconverged Infrastructure for Virtualization

## Features



Red Hat Hyperconverged Infrastructure (HCI) incorporates the latest Red Hat Virtualization technology along with Red Hat Gluster® Storage and Red Hat Ansibl® Automation. Red Hat Hyperconverged Infrastructure Virtualization is a collective group of proven technologies that provides a single support stack for virtual computes and virtual storage. Users can configure and deploy the components on the same Supermicro green technology server for a reduced IT physical footprint. This is achieved through Red Hat’s virtualization open software platform that virtualizes Linux and Microsoft Windows workloads leveraging Red Hat® Enterprise Linux® along with kernel-based virtual machine (KVM) technology that virtualizes resources, system processes and, customer applications. Red

Hat Gluster Storage is a software-defined and distributed scale-out file system that allows for rapid provisioning of storage based on the user’s consumption requirements. Red Hat Ansible® Automation is an enterprise automation platform that provides IT admins, with control and insight, managing policy enforcement and governance.

Features	Description
Choice of Supermicro Systems	<p>Supermicro offers systems with different form fit and function for different applications. Each server supports one or two Intel Scalable Xeon processors and up to 2+TB of system memory.</p> <ul style="list-style-type: none"> <li>• Ultra 2U offers high capacity storage, with each node supporting NVMe storage up to 384TB</li> <li>• BigTwin offers a very compact system that can host a 12-node RHHI-V cluster in 6 rack-unit</li> <li>• Other systems are supported as well</li> </ul>
Resource Saving / Green Computing	<p>Supermicro Resource Saving Systems that address the Global Environmental Impact of IT Infrastructure and provide improved TCO for Servers and All-Flash NVMe storage solutions.</p> <p>Benefits include:</p> <ul style="list-style-type: none"> <li>• Disaggregated Server Architecture with savings up to 45% to 60% in Hardware Refresh Costs</li> <li>• Free Air Cooling with up to 30% PUE Improvement</li> <li>• Pooled Resources with Dynamic Allocation and Configuration of Critical Datacenter Resources with up to 60% Savings in HVAC costs</li> <li>• Rack Scale Design with Open Industry Standard Rack Scale Management</li> </ul>
Simplified Plug & Play Scaling	Systems can start with 3 server nodes. As more performance and storage are needed, 3 server nodes can be added, up to a maximum of 12 server nodes.
Virtualization – Virtual Machine support	Red Hat Enterprise Linux with Kernel-based Virtual Machine (KVM) virtualizes resources, system processes, and customer applications. VMs running server workloads and virtual desktops are supported, including Windows and Linux virtual machines.
Replicated Storage	Red Hat Gluster Storage is a scalable, distributed file system that aggregates disk storage resources from multiple servers into a single global namespace. It is scalable to several PB, handles thousands of clients, POSIX compatible, accessible via NFS and SMB/CIFS. Resiliency features include replication, quotas, geo-replication, snapshots, and bitrot detection.
Virtual Data Optimization (VDO) module Deduplication and Compression	Virtual Data Optimizer (VDO) is a device mapper module that adds data reduction capabilities to the Linux block storage stack. VDO uses inline compression and data deduplication techniques to transparently shrink data as it is being written to storage media.
Ansible Automation	Ansible Automation utilizes hundreds of modules that automate IT environments and processes. It enables quick deployment of IT services, applications, and environments while automating routine activities. What traditionally took hours can now be accomplished in minutes. The automation performs system tasks, such as data replication to another system at a remote location.
Secure Virtualization or sVirt and Security-Enhanced Linus or SELinux	Secure Virtualization or sVirt and Security-Enhanced Linus or SELinux ensures security protection against hackers and viruses affecting the underlying system infrastructure. Supermicro supports TPMv2 to provide a hardware-based security infrastructure.
vGPU support	Nvidia vGPU support is an optional feature that accelerates graphics-intensive VDI workloads.

## Value of Supermicro Solutions for Red Hat Hyperconverged Infrastructure for Virtualization

Value	Features	Details
Save Time		
	Simple configurations and setup	After the simple installation of Red Hat OS and SSH security setup, RHHI-V takes 30 minutes to install.
	Simple scaling	User choice of 4 configurations: 3, 6, 9, or up to 12 nodes. The main choices for each server node is the number of cores in CPUs, amount of memory, and amount of storage.
	Automation such as replication	Red Hat Ansible Automation performs system tasks, such as data replication to another system at a remote location.
	Graphical management	Web based GUI interface.
Save Money		
	Cost effective systems	Supermicro systems including BigTwin, Ultra are cost-effective systems that support hundreds of TB per server.
	Single supported software stack	The purchase of the software subscription includes support. vGPU hardware sold separately.
	All RHHI-V software features are included	A single integrated environment, with Red Hat Enterprise Linux, LVM volume manager, Red Hat Virtualization, Red Hat Gluster Storage, and Red Hat Ansible Automation scripts to manage the environment.
	Choice of optimized systems for customer specific environments	Choice of systems to optimize racking, power, acquisition costs (Ultra 2U, BigTwin, other options).
	Reduced storage cost and network replication cost	Built-in deduplication and compression in the storage system. Deduplication significantly reduces the amount replicated data and time delivered to remote systems.
Worry-Free		
	Proven system design and quality	Supermicro has over 27-years of systems design experience, delivering over 10 generations of Intel systems, in a variety of form factors.
	Proven virtual environment to run Windows and Linux virtual machines	Red Hat Virtualization includes support for Windows and Linux virtual machines.
	Redundant data and reliable file system	Red Hat Gluster Storage keeps copies of all data in 3 servers. Red Hat LLVM volume manager has 20 years of proven reliability.
	Proven deduplication and compression	RHHI-V deduplication and compression are based on the reliable Virtual Data Optimization (VDO) module.
	Remote replication support	All data can be replicated to remote sites for disaster recovery.
	Integration with Red Hat infrastructure	RHHI-V is Red Hat integrates very well with any existing Red Hat infrastructure.
	vGPU graphic acceleration	Nvidia vGPU supports accelerated virtual desktop

Red Hat Hyperconverged Infrastructure can be deployed to complement your existing Red Hat Enterprise Linux infrastructure while seamlessly integrating with other Red Hat and Red Hat Partner products. RHHI-V management is built using Red Hat Ansible Automation scripts and playbooks. Red Hat Ansible automation is the foundation of the automation for your data center across organizations with integrated tools that allow for enterprise-wide automation. Ansible is an automation tool that configures and deploys software in your IT environment allowing you to orchestrate advanced tasks for continuous deployments or zero downtime while performing rolling updates. Ansible's goal is to add simplicity and ease of use while focusing on security reliability whether managing and deploying small environments or large enterprise deployments.

Data reduction is another valuable feature that allows IT organizations the capability of reducing data redundancy on large numbers of virtual machines. The integrated solution is part of the Virtual Data Optimization (VDO) module for the Linux device mappers provided within Red Hat Enterprise Linux. The incorporated Virtual Data Optimization module for Linux device-mapper provides inline, real-time and, block-level deduplication and compression at the 4KB level in conjunction with the management of the VDO data reduction. Deduplication reduces storage consumption through elimination of multiple copies of duplicate blocks. VDO takes the action of detecting duplicate blocks and recording them as a reference to the original block. After deduplication completion, multiple logical block addresses can then be mapped to the same physical block address that is defined as a shared block. When a shared block gets overwritten a new physical block is allocated for storing the new block data to ensure that other logical block addresses that are mapped to the shared physical block are not modified.

VDO also incorporates in, addition, to block-level deduplication inline block-level compression using the HIOPS Compression™ technology. The optimal solution is deduplication when running in virtual machine environments and backup applications where compression is more suited to structured and unstructured file formats that do not have block-level redundancy as with log files and databases. Compression works on blocks that have not been identified as duplicates. When unique data is discovered for the first time, it's in a compressed format. Any subsequent copies of data that have already been stored are deduplicated without any additional compression step. The compression feature is based on a parallelized packaging algorithm that enables it to handle multiple compression operations at once. After storing the blocks and responding to the requestor a best-fit packing algorithm finds multiple blocks when compressed fitting them into a single physical block. If it is determined that a particular physical block is unlikely to hold any additional compressed blocks, it gets written to storage where the uncompressed blocks are freed and reused.

Red Hat's vGPU support offers companies requiring complete graphics a way of delivering strong graphical performance in a virtualized environment. Through the support of vGPU's, the solution allows for increased visual clarity, improved solution performance while decreasing lag times to deliver virtualized graphics at remote site locations. The vGPU feature makes it possible to divide a physical GPU device into multiple virtual devices known as mediated devices. The mediated devices can then be assigned to multiple virtual machines (VMs) as virtual GPU's. The outcome is the VM's end up sharing the performance of a single physical GPU.

Open Network Virtualization or OVN is a software-defined networking feature which is an open-source virtual switching project that separates a physical network topology from the logical network. OVN improves scalability and performance while facilitating live migration of the virtual networking components without any hypervisor intervention. Virtual networking abstraction for OVS, implemented using L2 and L3 overlays that can also manage connectivity to physical networks. Support for flexible ACL's and native support for distributed L3 routing using OVS flows, with support for both IPv4 and IPv6. Additional support is available for ARP and IPv6 Neighbor Discovery suppression for known IP-MAC bindings and NAT and load balancing using OVS connection tracking. Native fully distributed support for DHCP also includes and works with OVS datapaths such as default Linux kernel datapath, DPDK or Hyper-V and includes software-based L2 gateways and TOR based L2 gateways that implement hardware\_vtep schema. And lastly, OVN enables support for networking on both VM's and containers running inside of the VM's without the second layer of overlay networking.

Security is one of the top IT concerns in the modern data center where RedHat has implemented Secure Virtualization or sVirt and Security-Enhanced Linux or SELinux to aid in the protection of the hypervisor against unknown threats targeted at the host or virtual machines. An added feature is the support of network encryption.

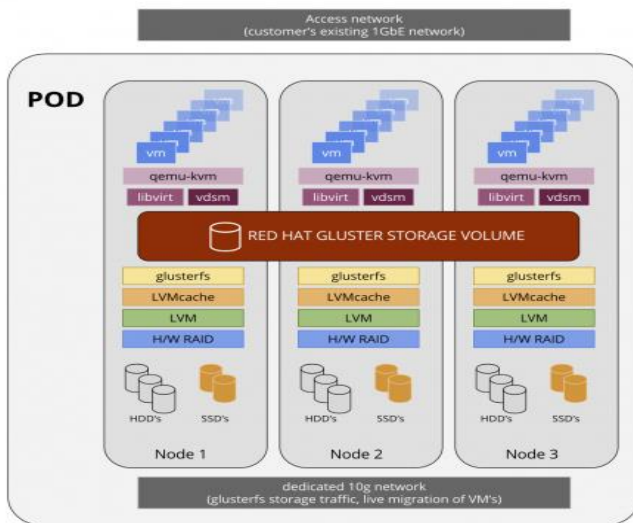
Central virtual resources management is a central management tool for compute, network, and storage through a web-based Red Hat Virtualization Manager encompassing a RESTful API that extends to management and

orchestration.

## Red Hat Hyperconverged Infrastructure Architecture Overview

Open Red Hat Hyperconverged Infrastructure (RHHI) combines compute, storage, networking, and computer management in one complete solution. With Red Hat Virtualization and Red Hat Gluster Storage RHHI deploys across three physical nodes to create a discrete cluster or pod. Remote Branch Offices (ROBO's) is a great use case for deployment with Supermicro servers. Essentially ROBO features remote office data synchronization to the central data center without having to rely fully upon to function.

The Single Cluster Architecture is highlighted.



Operating system support for Red Hat Hyperconverged Infrastructure (RHHI) is supported only with the Red Hat Virtualization host. Red Hat Hyperconverged infrastructure (RHHI) requires a minimum of 3 physical nodes while scaling to 6,9 or 12 nodes as supported solutions. For each physical machine, you must have at least 2 physical NICs per physical node. This is required for the separation of data and management traffic. Guidelines for deployments are as follows:

- Solutions: 12 cores, 64GB of RAM with 48TB of storage
- Medium Solutions: 12 cores, 128GB of RAM w64TB of storage
- Large Solutions: 16 cores, 256GB of RAM with 80TB of storage

To assist you with the sizing of a deployment you can refer to the Red Hat sizing chart that is located at the following URL: <https://rhhisize.com/>

Virtual Machines can have at most 4 virtual CPUs and 2TB virtual disk space.

Networking – Each node requires 3x1 Gigabit Ethernet ports. For high availability, both ports must split across two network switches with separate power supplies for fault tolerance. Front end networking management that is used by Red Hat Virtualized and virtual machines should be running at minimum Gigabit Ethernet speeds. All IP addresses assigned to the front-end network must be on the same subnet as each other. All IP addresses on this network must 'not' be on the same subnet as the back-end storage and migration network.

Back-end network for the storage and migration traffic between storage peers; the recommendation is to have a 10Gps network with Red Hat Gluster Storage requiring a maximum latency of 5 milliseconds between peers.

For IPMI fencing, a separate network is required. For DHCP network configurations for the Host Engine virtual machine you must have a DHCP configured before the configuration of the Red Hat Hyperconverged solution. To allow geo-replication to store copies of data for disaster recovery purposes, a reliable time source is required.

Storage - If it's top performance you're looking for in relation to disk technology then the recommendation is to use SSD's in place of Hard Disk Drives. RAID configurations are dependent upon the chosen technology.

- SAS/SATA 7K disks are supported with RAID 6
- SAS 10K and 15K disks are supported with RAID 5 and RAID 6 configurations
- RAID cards must use flash backed write cache
- For JBOD configurations you'll want to contact a sales representative

Logical volumes that comprise the engine gluster volume must be thick provisioned. This protects the Hosted Engine from out of space conditions, disruptive volume configuration changes, I/O overhead, and migration activity. The logical volumes that comprise the vmstore and optional data gluster volumes must be thin provisioned. This allows for greater flexibility in underlying volume configuration. If your thin provisioned volumes are on Hard Drive Disks (HDD's), you'll want to configure a smaller, faster Disk (SSD) as a lvmcache.

Red Hat Gluster Storage Volumes in a RHHI-V configuration are expected to have 3 to 4 in a solution.

- 1 engine volume for the Hosted Engine
- 1 vmstore volume for virtual machine boot disk images
- 1 optional data volume for other virtual machine disk images
- 1 shared\_storage volume for geo-replication metadata

Also, a Red Hat Hyperconverged infrastructure deployment can contain at most one geo-replicated volume.

Volume Types - Red Hat Hyperconverged Infrastructure (RHHI) supports only replicated or arbitrated replicated volume types. The arbitrated replicated volume type carries the following additional limitations.

- Arbitrated replicated volumes are supported on the first 3 nodes only
- Arbitrated replicated volumes must use replica 3 arbiter 1 configuration
- Arbiter bricks do not store file data; they only store file names, structure, and metadata
  - For a three way-way arbitrated replicated volume requires about 75% of the storage space that a three-way replicated volume would require to achieve the same level of consistency.
  - The arbiter brick stores only metadata, a three-way arbitrated replicated volume only provides the availability of a two-way replicated volume

Support Limitations – One arbitrated replicated volume is supported as part of the initial deployment of Red Hat Hyperconverged Infrastructure (RHHI). Expanding an arbitrated replicated volume is not supported. Adding additional arbitrated replicated volumes is not supported. And further, only arbitrated replicated volumes are currently supported only in the following configuration, replica 3, arbiter 1.

Disaster Recovery – Red Hat strongly recommends configuring a disaster recovery solution to synchronize data from one Red Hat Gluster Storage cluster to another. By synchronizing the local data volume from your discrete Red Hat Hyperconverged Infrastructure (RHHI) cluster to a central data center on a regularly helps ensure you can restore your cluster to a working state after an outage.


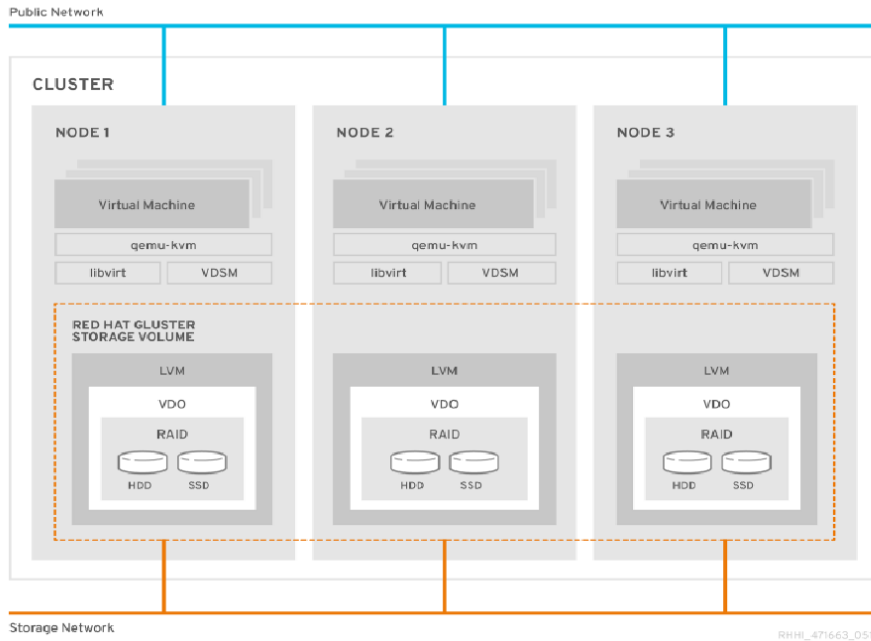
# Supermicro Solutions and Reference Architectures for Red Hat RHHI-V

## Supermicro Red Hat HCI Solutions



	Balanced	Capacity	Capacity Performance	Density Performance
Solution SKU				
System SKU	SYS-6029U-E1CRT-1-RHHI	SYS-2029U-E1CR4-1-RHHI	SYS-1029U-TN10RT-RHHI	SYS-2029BT-HNR-01-RHHI
Cluster Size	3 Nodes	6 Nodes	12 Nodes	12 Nodes
Compute Cores	96	240	672	672
Total Memory	384GB	1152GB	4608GB	4608GB
Total Storage (Raw)	72TB (2.0TB SAS3)	138TB (3.84TB SATA3)	230TB (1.92TB NVMe)	138TB (1.92TB NVMe)
Networking	2 x 10GBASE-T or 2 x 10G SFP+		2 x 25G SFP28	
Software	Red Hat Virtualization (RHVH) 4.3			

## Red Hat Hyperconverged Infrastructure for Virtualization Architecture



Supermicro RHHI-V Medium Configuration



# Supermicro BigTwin: The Industry's Highest Performing Twin Multi-Node System

First and only 2U Multi-Node Systems supporting 205-watt dual-Xeon Processors, 24 DIMMs per node, and 24 All-Flash NVMe



- No-compromise 5th Generation Twin architecture delivers the highest performance and efficiency in a 2U 4-node platform
- The first and only multi-node system that supports the widest TDP range of CPUs (up to 205-watts), fully exploits all memory channels with 24 DIMMs per node and 24 All-Flash NVMe drives
- 30% better thermal capacity in a compact 2U form factor enables configurations with the highest performing processor, memory, storage and I/O leveraging the Titanium level 96% efficient PowerStick power supplies (2200W/2600W)
- 100% more I/O capacity and added flexibility with more than 10 networking options including 1GbE, 10G, 25G, 100G, IB and industry leading SIOM modular interconnect
- Each node can support current and next generation dual Intel Xeon processors with up to 6TB of memory, 24 drives of All-Flash NVMe, Hybrid NVMe/SATA/SAS, SSD and HDD, plus 2 m.2 NVMe/SATA per node

## THE 5TH GENERATION TWIN ARCHITECTURE


The Supermicro® BigTwin™ is a breakthrough modular multi-node server system that eliminates traditional modular vs. rack computing design trade-offs. BigTwin is the 5th generation in the patented Supermicro Twin Family with a multitude of innovations and engineering breakthroughs. The no-compromise architecture delivers the efficiency of multi-node density optimized systems with the full features and power of traditional 1U and 2U rack mount systems.



Figure 1. Rear View of 2U BigTwin Server System Showing 4 Hot-swap DP Nodes and 2 Redundant Power Supplies

## A No-Compromise Design

The key benefit of BigTwin is its no-compromise design. Historically multi-node systems traded off features and capacity for higher density. They were deployed for workloads that did not require the highest performance or the highest memory density on a single node. The new 2U BigTwin design is a breakthrough multi-node system that supports the highest performing CPUs, full 24 DIMMs of memory and up to 24 all-flash NVMe SSD drives.



24x 2.5" Hot-Swap Drive Bays Supporting NVMe/SAS/SATA drives (6 driver per node)

Figure 2. Front View of 2U BigTwin Server System (All-Flash NVMe Model)

## ALL-FLASH NVME AND 24 DIMMS PER NODE

SIOM Modules	Networking Features
AOC-MH25G-m2S2T	Dual-port 25 Gigabit Ethernet SFP28 and Dual-port 10 Gigabit Ethernet RJ45
AOC-MHIBF-m2Q2G	Dual-port FDR InfiniBand QSFP+ and Dual-port Gigabit Ethernet RJ45
AOC-MTG-i4S	Quad-port 10 Gigabit Ethernet SFP+
AOC-MGP-i4	Quad-port Gigabit Ethernet RJ45

Supermicro has been a leader in introducing NVMe support and advanced NVMe features in our product portfolio and BigTwin continues that tradition. BigTwin supports 24 hot-swappable 2.5" U.2 NVMe drives, or mixed configurations with SAS3 and SATA3 drives (varies by different BigTwin models).

Supermicro's, BigTwin utilizes a shared design and common components; most multi-node systems are one-off designs that limit the rest of the product family in regards to providing sub-optimal server, sparing, and training. Having a building block design allows the user to optimize a specifically intended workload.

The 24 DIMMs can deliver better performance and responsiveness for database and cloud infrastructure workloads, as well as providing the flexibility to customers looking to invest in lower capacity DIMMs for the lowered cost of ownership.

## Applications

A hyper-converged solution of 3 nodes plus a hot spare can be hosted by just one BigTwin server. Each node can support dual Intel Xeon E5-2600 v4/v3 processors with up to 3TB of memory, 2 NVMe SSD drives for caching, and 4 SAS drives for datastore.

The Supermicro BigTwin is a breakthrough modular multi-node server system, an ultra-dense an energy-efficient powerhouse that provides industry-leading performance-per-watt, per-dollar, and per-square-foot.

## 1GB/S TO 100GB/S

Supermicro Super I/O Module (SIOM), common across both Rack and Modular platforms, delivers up to 50% I/O cost savings and freedom to select networking options from 1Gb/s to 100Gb/s. The SIOM enables a tighter system integration and leaves a total of two PCI-E 3.0 x16 LP slots for storage or networking add-on-cards on each node.

## Supermicro Ultra SuperServer Delivers Unrivaled Performance for Enterprise Workloads

Ultra SuperServer is designed to deliver unrivaled performance, flexibility, scalability, and serviceability for demanding workloads



- Supports Dual Intel Xeon Scalable processors with up to 205W TDP
- 1U and 2U Rack-mountable with up to 32 DIMM slots with 4TB DDR4 memory, Intel Optane DC persistent memory
- Up to 24 U.2 NVMe drives, with 2.5" and 3.5" SAS3, SATA, and M.2 options
- Up to 8 PCI-E slots with flexible onboard 1/10/25G Ethernet options
- Compact form factor enables configuration with leveraging 1600W / 1000W / 750W Redundant Titanium (96%) and Platinum Level Power Efficiencies
- Management tools with Open Industry standard IPMI, Redfish API's, Rack Scale Management coupled with World Wide Global Services Support

## Ultra-Servers for the Highest Performance and Flexibility for Enterprise Applications

Supermicro Ultra SuperServers are designed to deliver the highest performance, flexibility, and serviceability to demanding IT environments and to power mission-critical Enterprise workloads, including support for the 2<sup>nd</sup> generation Intel Xeon Scalable process and a new class of big, affordable memory – Intel Optane DC Persistent Memory. The Intel UltraPath Interconnect (UPI) and Crossbar technologies provide direct communication between two CPUs in a dual-socket server to catapult performance and dramatically reduce latency.

Available in 1U and 2U form factors, Ultra SuperServers support up to 6TB DDR4-2933MHz memory in a 24 DIMM configuration with SATA3 and optional SAS3 and NVMe support for increased storage bandwidth that includes a variety of Ultra Riser options, including built-in 1G, 10GBASE-T, 10G SFP+ and 25G Ethernet options with support for add-on SAS3 HW/SW RAID controllers and additional PCI-E 3.0 slots. The perfect range for diverse workloads and applications and can be easily reconfigured for multiple Enterprise and Data Center applications in Virtualization, Big Data, Analytics, and Cloud Computing.



## Industry Applications for Supermicro and Red Hat RHHI-V HCI Solutions

### Hyper-Converged Infrastructure Remote Office Branch Office (ROBO) Solutions

As businesses adjust to changing economic and physical environments they continue to search for optimized growth and ways to increase productivity while lowering cost. Medium to small companies are looking to improve Remote Office / Branch Office (ROBOs) as one of their top priorities. Through IT optimization ROBO's have the capacity to improve employee productivity, customer retention while improving the total cost of ownership.

Red Hat HCI ROBO solutions offer the power of Supermicro Servers while enhancing the capabilities of enterprise-class deployments while offering increased performance, data protection, data efficiencies and a robust unified management tools to simplify the complexity of management. Red Hat RHH-V HCI allows for the deduplication of data while offering compression will optimize data throughout the product life cycle on a global basis. All Supermicro server CPUs are reserved for the production of VM's rather than separate infrastructure applications. ROBO offers improved HCI data protection while eliminating the requirement for backup and recovery technology required with traditional converged infrastructures. Red Hat HCI can be managed globally and independently regardless of the location where you have access to managing all ROBO's.

Red Hat RHHI-V provides a lower cost of ownership through optimized ROBO environments through lower prices and fewer nodes for full availability. HCI management is simplified through ROBO IT with the integration of IT products below the hypervisor and by combining a globally unified management system across all remote sites. Red Hat ROBO solutions offer better application performance, faster access with greater data protection while offering enhanced WAN effectiveness and efficiency, and global management across ROBO environments.


## HCI Hardware Choices Matter for Critical HCI Solutions

We often put heavy emphasis on the software in the software-defined HCI infrastructure while never fully evaluating the hardware and the features that differentiate it from the pack. The goal of having a high performing power-oriented HCI solution is to have the software and hardware working as a most optimal HCI package. Supermicro provides a Red Hat HCI turn-key solution, integrated appliance in a rack-based solution that's tested and optimized for the Balanced, Capacity, Capacity Performance and Density Performance workloads. By offering a turnkey solution for your line of business that is pre-configured and tested, the predictability provides enhanced value.

## Suggested Applications

Supermicro servers combined with RHHI-V offer many possibilities to run IT systems better. The best environment to deploy RHHI-V based solutions are as follows:

- < up to 1000 users
- < 1PB storage
- Want data redundancy and resiliency
- Possibly need data replication to multiple sites or to main data center
- Windows or Linux applications, for servers and for desktop
- Customers have working experience with Red Hat infrastructure or willing to switch to Red Hat based infrastructure



	Server Workloads	Virtual Desktop Workloads
Smaller IT needs	Windows or Linux server applications: <ul style="list-style-type: none"> <li>• Email</li> <li>• File &amp; Print</li> <li>• Windows or Linux databases, analytics</li> <li>• Web applications</li> </ul>	Windows or Linux desktop applications: <ul style="list-style-type: none"> <li>• Windows 10 desktop</li> <li>• Linux desktop</li> <li>• CAD / CAM / Simulation</li> <li>• Media editing</li> <li>• Recommend 1 core / 2 thread per virtual desktop</li> <li>• For Windows 10 desktop and graphics intensive applications, recommend using GPU with vGPU support. GPU could be NVIDIA T4, RTX6000, RTX8000, or V100</li> <li>• NVMe storage is recommended if high end graphics applications are used.</li> </ul>
Remote Office / Edge (Retail, Banking, Telecommunications,	Windows or Linux server applications: <ul style="list-style-type: none"> <li>• Email</li> <li>• File &amp; Print</li> </ul>	

Public Sector,  
Energy)

- Industry specific custom applications
- NVMe storage is recommended if high end graphics applications are deployed
- Size the number of cores needed to run server applications
- Size memory accordingly

Considerations:

1. Rack space available to house the servers:  
For a total less than 6U, use BigTwins
2. Size the number processor cores that needs to run all the applications. Divide this total number cores into the number of servers, divided by 2 (to account for 2 processors in each server), to select the processor needed. Also, note the CPU speed needed.
3. Size memory needed by the applications.
4. Size storage
5. Select the network speed needed: 10GbE, 25GbE, or 100GbE.

## Summary

Supermicro systems combined with Red Hat Hyperconverged Infrastructure for Virtualization (RHHI-V) offers a simple, low-cost, and proven hyperconverged infrastructure (HCI) with the right building blocks to run applications and data storage for edge and branch offices. The systems can easily scale from a small, medium to large configuration all dependent upon your application and workload requirements.

## Resources

For more information, please contact your Supermicro representative.

Supermicro Systems with RHHI-V:

[https://www.supermicro.com/solutions/Solution-Brief\\_RedHat\\_Open\\_Hyperconverged.pdf](https://www.supermicro.com/solutions/Solution-Brief_RedHat_Open_Hyperconverged.pdf)

Supermicro Servers:

<https://www.supermicro.com>

Red Hat Hyperconverged Infrastructure for Virtualization:

<https://www.redhat.com/en/technologies/storage/hyperconverged-infrastructure>