



TABLE OF CONTENTS

2 INTRODUCTION

3 HARDWARE AND SOFTWARE CONFIGURATION

- Hardware Component Overview
- Component Firmware
- Management Software Versions
- Network Topology
- PCI-E Topology

7 INSTALLATION

- Installation Overview
- Switch Configuration
- RSD Installation
- Configure RSD
- Composing Nodes with Pooled NVMe Drives

18 APPENDIX

- Supermicro BigTwin Servers
- NVMe Enclosure Technical Details
- Mounting NVMe Drives
- Partitioning NVMe Drives
- Dynamic PCI-E Attachment/Detachment
- Reference BOM

WHITE PAPER


SUPERMICRO RSD HIGH PERFORMANCE LARGE SCALE NVMe STORAGE REFERENCE DESIGN

EXECUTIVE SUMMARY

This document describes a Reference Design for large scale, high performance NVMe storage solutions built from Supermicro Server and Storage product offerings and managed by Supermicro Rack Scale Design software. The design enables simplified and rapid deployment, management, and expansion of petabyte scale NVMe storage platforms. These platforms can provide the foundation for scaled out High Performance Compute local storage or cloud scale Software Defined Storage deployments.



**SUPERMICRO 2U 4-NODE
BIGTWIN™ SUPERSERVER**



**SUPERMICRO 1U 32 U.2 NVMe
STORAGE ENCLOSURE**

Super Micro Computer, Inc.
980 Rock Avenue
San Jose, CA 95131 USA
www.supermicro.com

INTRODUCTION

NVMe based storage is rapidly becoming the standard in datacenters for high speed storage applications. The ability to rapidly and reliably deploy and operate scaled NVMe storage systems is critical to maximizing the benefit offered by this new technology. A complete integrated offering of hardware and software components can greatly simplify the lifecycle for advanced flash storage solutions. Supermicro is able to offer such an offering through a combination of state of the art NVMe optimized BigTwin servers, 1U 32 NVMe storage enclosures and Supermicro Rack Scale Design software.

This document describes a Reference Design for an extreme performance NVMe storage solution based on a building block model allowing easy scaling to multiple petabytes of flash. This design can be used as the foundation for High Performance Computing with large scale local flash or in Software Defined Storage systems providing network attached storage for scale out private cloud workloads.

The Reference Design is composed of a number of subsystems, each providing a specific function of the solution. Storage is provided by Storage Blocks containing servers, storage enclosures, and NVMe drives. Storage blocks are deployed as needed to provide the total capacity and performance required by the target workloads. Figure 1 shows the rear view of a single 3U Storage Block. Up to 12 Storage Blocks can be supported in a single 42U rack.



Figure 1. *NVMe Storage Block (rear view)*

Supermicro Rack Scale Design software is used to deploy and manage the Storage Blocks at rack level granularity. With Supermicro RSD racks can be populated with NVMe flash Storage Blocks, dedicated Compute Nodes, and other forms of storage, along with the required networking in the exact ratios required by target workloads. Resource management and footprint expansion are greatly simplified with the RSD POD Manager interfaces and can be readily integrated with datacenter orchestration tools through the POD Manager Redfish Restful APIs. The use of Template based management in Supermicro RSD allows newly added resources to be Composed and Assembled in only a few steps.

HARDWARE AND SOFTWARE CONFIGURATION

HARDWARE COMPONENT OVERVIEW

Management Subsystem

Table 1 shows the components providing the Management Subsystem for the Reference Design. A single RSD Appliance is used to manage multiple racks of Storage Blocks. Each rack includes local Management and Data switches. The Management Switches are linked to provide a complete Management Network for the RSD appliance to manage the various subsystem components.

Table 1. *Management Subsystem Components*

PARTS	DESCRIPTION
RSD Appliance and Software	SYS-5019S-TN4-SRSMGT and SFT-SMCI-SRSDM
Management Switch	Supermicro SSE-G2252
Data Switch	Supermicro SSE-C3632SR 32 port 100Gb As required for production traffic

Server Subsystem

The server subsystem, shown in Table 2 is composed of high performance, NVMe optimized Supermicro BigTwin servers. BigTwin servers offer high density (four servers per 2U) and maximum performance through the complete range of Xeon Scalable Processors, maximum memory speed and capacity, NUMA optimized network and storage placement, and support for large amounts of NVMe storage. The Server Subsystem is connected to the production network through multiple 100Gb RDMA capable interfaces.

Table 2. *Server Subsystem Components*

PARTS	DESCRIPTION
BigTwin NVMe hosts (4)	SYS-2029BT-HNR
PCI-E Expansion Card	AOC-SLG3-4X4P-P
Network (2)	Mellanox ConnectX-4 PCI-E

NVMe Flash Enclosure Subsystem

The NVMe Flash Enclosure subsystem, shown in Table 3, provides for the attachment of additional NVMe drives to the servers in a BigTwin chassis. The combination of NVMe capacity local to the BigTwin chassis and in the NVMe Enclosure provides the extreme, high density capacity achieved in the Reference Design. The NVMe enclosure adds an additional 32 U.2 NVMe drives to the attached BigTwin servers.

Table 3. *NVMe Flash Enclosure Subsystem Components*

PARTS	DESCRIPTION
NVMe JBOF enclosure	SSG-136R-N32JBF
PCI-E x8 Cables	CBL-SAST-1035-1

FOR MORE INFORMATION

- Supermicro® MicroBlade™ Solutions
www.supermicro.com/products/MicroBlade/
- Supermicro Rack Scale Design
www.supermicro.com/solutions/SRSD.cfm
- Intel® Xeon® Processor
www.intel.com

NVMe Storage

The flash storage used in the BigTwin chassis and the NVMe enclosure is provided by standard U.2 form factor drives common to each subsystem. The Reference Design uses Intel P4500 4TB U.2 NVMe drives. Other drive models and capacities are compatible with this design.

Table 4. *NVMe Storage*

PARTS	DESCRIPTION
NVMe Drives	HDS-2VD-SSDPE2KX040T701 Intel P4500 4TB U.2 NVMe

COMPONENT FIRMWARE

Table 5 shows the firmware versions used during the validation of this design. Newer firmware may be required. Please follow current recommendations from Supermicro.

Table 5. *Firmware Versions*

PARTS	DESCRIPTION
BigTwin Firmware	9.12
BigTwin BIOS	T201710271459
NVMe Chassis Firmware	1.10
Management Switch Firmware	2.0.0.11
Data Switch Firmware	Cumulus Linux 3.4.2

MANAGEMENT SOFTWARE VERSIONS

Table 6. *Software Versions*

PARTS	DESCRIPTION
Supermicro RSD	2.1.3 build 171110

NETWORK TOPOLOGY

In the Reference Design, the RSD Management Network is private and air gapped from other data or management networks. Supermicro RSD provides addresses to managed components on the Management Network from a private DHCP server. Supermicro RSD supports uplinking of the Management Network to customer Management Networks when direct access to managed entities, such as server BMCs is desired, but is outside the scope of this document.

The Data Network is uplinked as necessary to support the use of the Storage Blocks to datacenter applications. It is supported to be kept private to the RSD POD.

Management Switch

Table 3-7 shows the port connections for the Management Switch. If the Data Switch provides a dedicated management port, it should be used for the connection to Port 44 of the Management Switch, otherwise an available port on the Data Switch should be used and configured to be the Management Port.

Table 7. *Network Addresses*

PORT	CONNECTION
1-42	Server Node BMC
44	Data Switch Management
45	Previous Rack Management (Port 46)
46	Next Rack Management (Port 45)
48	RSD Appliance LAN3

Data Switch

The Data Switch should have a connection to the Management Switch as described for Management Switch connections. The Data Switch may also be optionally uplinked to the customer Data Network if external access to RSD Compute resources is required. The LAN ports of the RSD Compute Nodes should be connected to the Data Switch.

RSD Appliance Connections

The LAN connections to the RSD Appliance are shown in Table 8. The LAN1 connection provides access to the Supermicro RSD POD Manager. It should be connected to the customer network used for system management and orchestration. LAN2 is the appliance IPMI BMC connection used to manage the appliance (including access to the console.) LAN3 provides the connection to the RSD private management network.

Table 8. *RSD Appliance Connections*

INTERFACE	CONNECTION
LAN1	Public Network (used to connect to POD Manager)
LAN2	BMC (Used to manage appliance, including installation)
LAN3	Management Switch Port 48

PCI-E TOPOLOGY

The BigTwin servers provide PCI-E lanes to internal chassis NVMe slots and to PCI-E slots available for Add-on Cards. The Reference Design connects the NVMe enclosure to the BigTwin servers using cabled PCI-E connectivity via a PCI-E Expansion Add-on Card.

This provides a total of 40 lanes per server of PCI-E for NVMe drive attachment. There are 24 lanes on CPU 1 and 16 lanes on CPU 2 (per server). Up to six drives per server in the BigTwin chassis and up to four drives per server in the NVMe chassis can be connected without PCI-E lane oversubscription. An additional four drives per server can be attached in the NVMe chassis for additional capacity.

See Section 5.2 for further information on connectivity and performance considerations for the NVMe Enclosure.

INSTALLATION

In this Installation example a single Storage Block is installed and configured. Additional Storage Blocks are easily added by extending the sequence for additional BigTwin Chassis and NVMe enclosures. The template based allocation model used in this design allows for the rapid provisioning of newly added racks of Storage Blocks.

INSTALLATION OVERVIEW

1. Configure Network Switches (once per rack)
2. Install Supermicro RSD on the RSD Appliance (once per POD)
3. Configure Rack into Supermicro RSD (once per rack)
4. Compose and Assemble Storage Blocks (as needed)

SWITCH CONFIGURATION

Each Management Switch needs VLAN 4050, 4053, and 4054 tagged on port 47 and port 48.

The Management Switch management IP should be configured. In this example it is set to 192.168.1.1. Additional switches can be configured similarly as needed.

The Data Switch management IP addresses should likewise be configured. In this example the Data Switch IP is set to 192.168.1.2.

SNMP should be enabled on all switches. Make a note of the configured Community Strings.

RSD INSTALLATION

Installing RSD

Pre-requisites

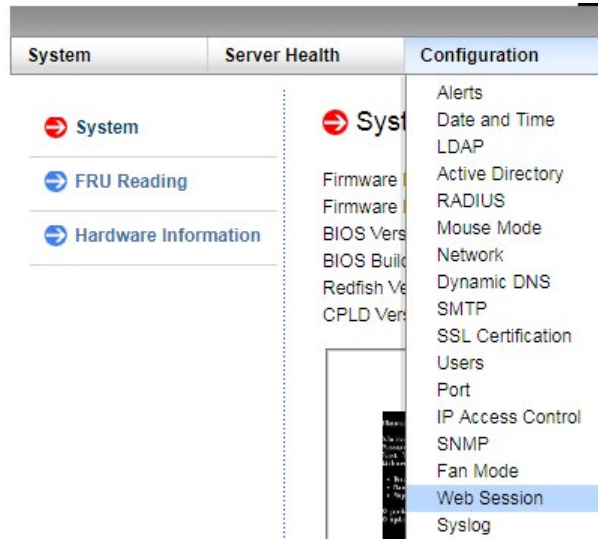
- Supermicro RSD installation ISO
- IP Addresses for the RSD Appliance BMC, public host interface, and POD Manager

These can be statically assigned or supplied by a datacenter DHCP server

Set and/or obtain the BMC IP address of the RSD Appliance by connecting a VGA cable to the monitor port.

Connect to the IPMI interface using a web browser. Login using default ADMIN/ADMIN credentials.

In Configuration > Web Session, enter 0, then Save.




➔ Web Session

This page allows you to configure web session parameters.

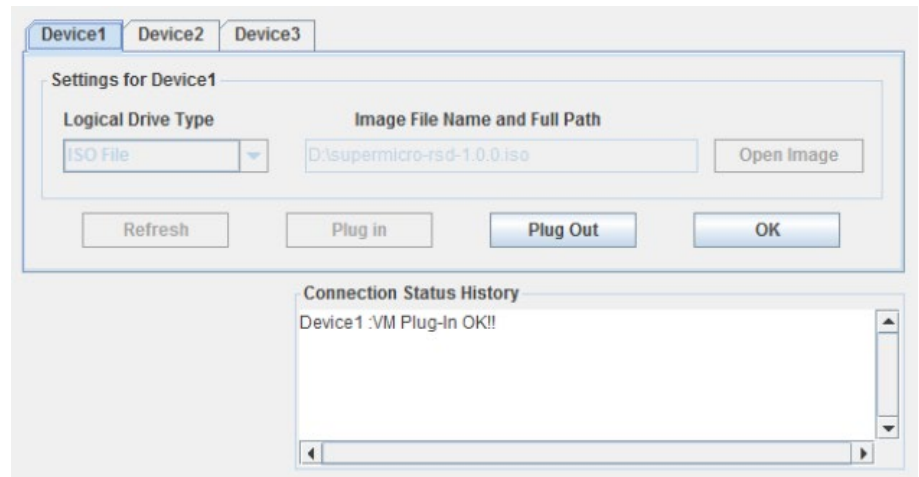
Session Timeout Value (minutes, range: 1-30, 0: never timeout)

Select Console Redirection (Java version).



Click on Launch Console. A jnlp file will be downloaded. Open it using Java Web Start.

Mount the Supermicro RSD ISO in Virtual Storage.

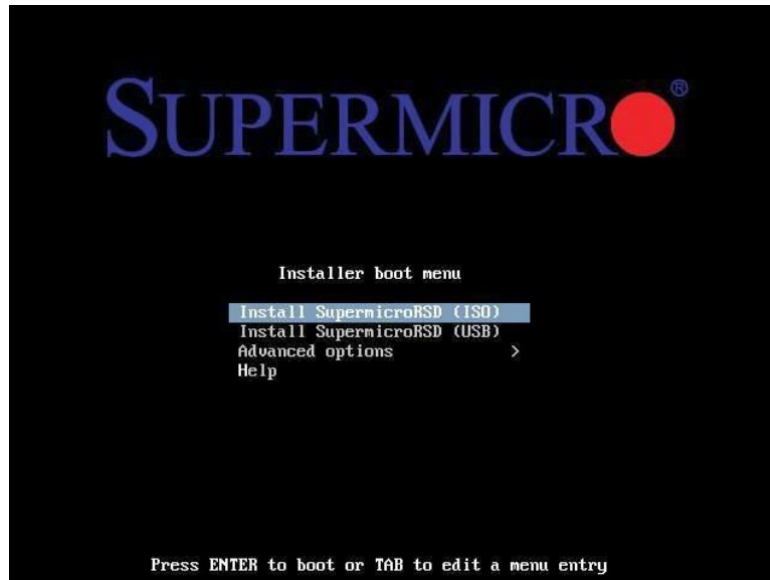


Boot the appliance by pressing F11.

Select ATEN Virtual CDROM YSOJ.



Select Install SupermicroRSD (ISO).



Configure static or DHCP IP address for the RSD appliance and Pod Manager.

When complete a login prompt will be displayed.

Login and attempt to ping the configured switch IPs.

CONFIGURE RSD

Open a browser window to the Supermicro RSD Drawer Configuration Editor:

<http://<RSD Appliance IP address>:8888/rsd/drawerconfig>

Fill in the First Management Switch information, including the SNMP community string.

First Management Switch

Switch IP Address

192.168.1.1

SKU

SSE-G2252

SNMP Community

Location

2

Height (U)

1

Fill in the RSD Appliance information.

1. Click "+ Rack"

RACK LIST + Rack Last Rack All

Click "+ Drawer"

Choose "Twin Family" from the drop down menu.

Enter the drawer location as installed in the rack and enter "2" for drawer height.

For each node, Click "+ Node", then enter the locations/BMC switch port numbers.

Server A is location 1, B is 2, etc.

Check the "Use Firmware Extensions for Deep Discovery" checkbox.

DRAWER-9

DRAWER-10

DRAWER-10 Twin Family

Drawer Location
10

Drawer Height (U)
2

SKU

☒ Use Firmware Extensions for Deep Discovery (X11 and later servers only)
☐ FatTwin Series

System Nodes (Modules)

Node Location	Node's BMC connects to Mgmt. Switch Port#	JBOF Cable IDs (test)	
1	3		✖
2	4		✖
3	7		✖
4	8		✖

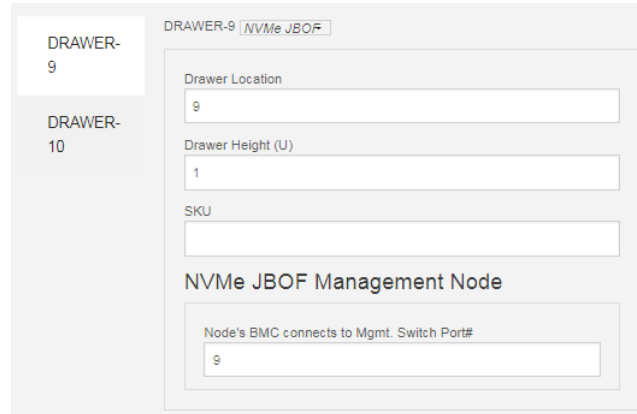
+ Node
Last Node
All

Click "+ Drawer"

Choose NVMe JBOF.

Enter the drawer location as installed in the rack.

Enter the BMC switch port number.



Drawer-9 | NVMe JBOF

Drawer Location: 9

Drawer Height (U): 1

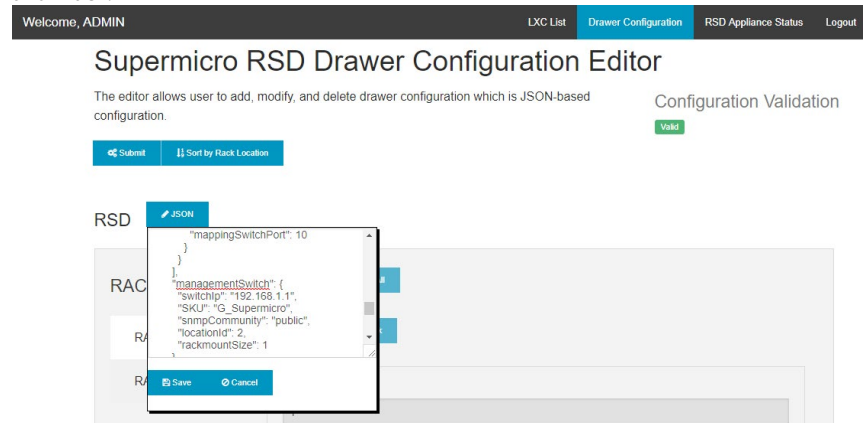
SKU:

NVMe JBOF Management Node

Node's BMC connects to Mgmt. Switch Port#: 9

2. Repeat from Step 1 for additional Storage Blocks.

For large additions it is possible to edit the raw JSON drawer configuration object by clicking the "JSON" edit button, modify the JSON data, and then click Save. The text in the pulldown text box can be selected, copied to an external editor and then copied back to the text box for more convenient editing. Blocks of text representing existing drawers can be duplicated, modifying the unique properties, to quickly generate the configuration for a full rack.



Welcome, ADMIN LXC List Drawer Configuration RSD Appliance Status Logout

Supermicro RSD Drawer Configuration Editor

The editor allows user to add, modify, and delete drawer configuration which is JSON-based configuration.

Configuration Validation: Valid

Submit Sort by Rack Location

RSD JSON

```
{
  "mappingSwitchPort": 10
},
{
  "managementSwitch": {
    "switchip": "192.168.1.1",
    "SKU": "G_Supermicro",
    "snmpCommunity": "public",
    "locationId": 2,
    "rackmountSize": 1
  }
}
```

Save Cancel

Supermicro RSD Drawer Configuration Editor

The editor allows user to add, modify, and delete drawer configuration which is JSON-based configuration.

Submit Sort by Drawer Location

The RSD POD Manager should now discover the nodes and NVMe Enclosure. Go to <http://<RSD appliance IP address>:8888/rsd/lxc> to see the status. It should show something like

Drawer Location	rest-server	chassis-agent	compute-agent	compute-ex-agent	network-agent	compute slots
4 FW Extension	on UUID	on v2.1.3.0.0	--	on v2.1.9.1.0	--	[1]:192.168.5.8 [3]:192.168.5.1 [2]:192.168.5.3 [4]:192.168.5.2

for the BigTwins, and something like

Drawer Location	rest-server	pnc-agent	JBOF Node
6	on UUID	on v2.1.3.0.3	192.168.5.10

for the NVMe Enclosures.

Go to <https://<Pod Manager IP address>:8443/smcipodm/> and login using ADMIN/ADMIN credentials.

After logging in, you will see the Dashboard page. Go to Physical Assets > Assets. Nodes should show up with their locations identified by Drawer location + Node number. Discovery may take approximately 10 minutes to complete or more for larger configurations.

COMPOSING NODES WITH POOLED NVME DRIVES

For this design RSD Templates are used for rapid composition of Storage Nodes from available Storage Blocks. Templates greatly simplify the steps to assemble and deploy newly added racks.

The following example template will compose a node from any Compute Node with an attached NVMe enclosure that has at least four available NVMe drives. Additional selection criteria (such as server processor parameters, memory capacity, or NVMe drive capacity) can be added for heterogenous deployments where a variety of resources are available for composition.

```
{
  "Name": "PNC-4",
  "LocalDrives": [
    {
      "Type": "SSD",
      "FabricSwitch": true
    },
    {
      "Type": "SSD",
      "FabricSwitch": true
    },
    {
      "Type": "SSD",
      "FabricSwitch": true
    }
  ]
}
```


```

    },
    {
        "Type": "SSD",
        "FabricSwitch": true
    }
]
}

```

Create the template

1. Login to the Supermicro RSD Web UI at <https://<POD Manager IP>:8443/smcipodm>.
2. Select "RSD Composed Node -> Composed Template Management" from the menu bar.



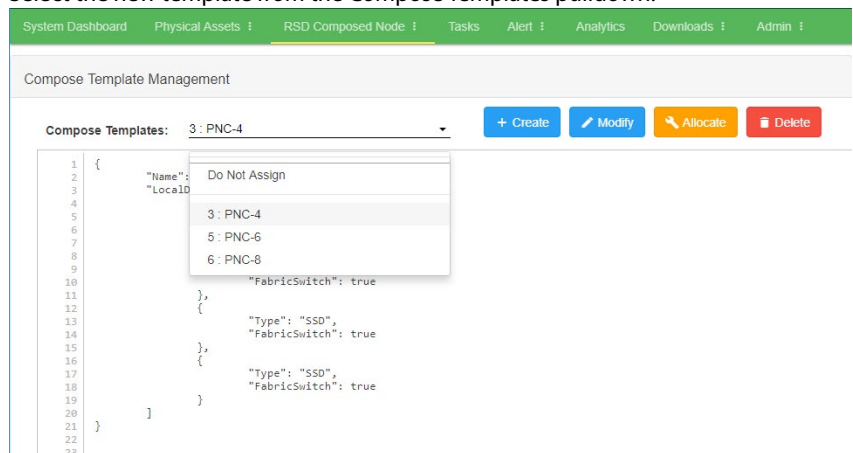
3. Copy and paste the template text into the numbered text box area, modifying as needed for local purposes.
4. Click "+ Create" and provide a name for the new template.

Note: The name in the template is for the eventual composed node and is not the template name.

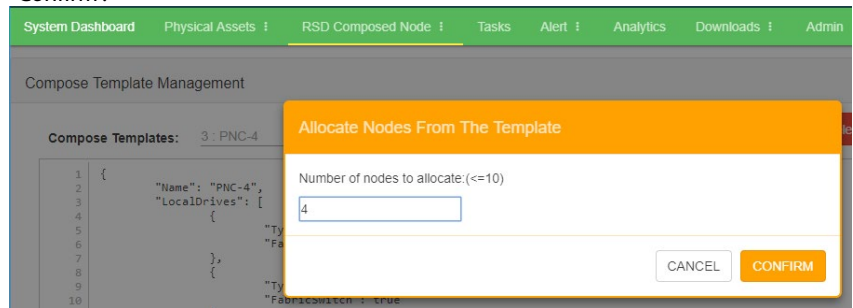
Compose nodes using the template

The newly created template makes it trivial to compose a large number of nodes from available Storage Blocks in a single operation.

1. Select the new template from the Compose Templates pulldown.




- Click "Allocate" and enter the desired number of Composed Nodes and click "Confirm".

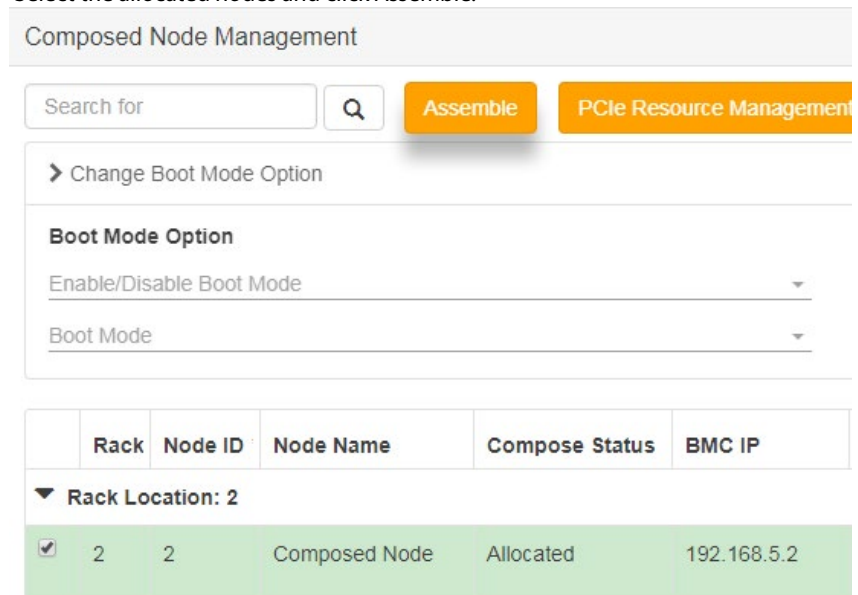


Assemble Composed Nodes

- Go to RSD Composed Node > Composed Node Management.



- Select the allocated nodes and click Assemble.



View allocated PCI-E resources

Go to Physical Assets > Assets, where Fabric will show the physical topology of each NVMe Enclosure.

Compute

Network I

Storage

Chassis

Fabric

Fabric:

14-ft-1

Refresh

Fabric Summary

0

1

Node ID: 2, System ID: 15-s-13, BMC IP: 192.168.5.2

Up

Zone Count: 4

Endpoint Count: 23 SSD/4 RootComplex

Chassis Type: Enclosure


Chassis Part Number: CES-F424AS-R1K28BP

Chassis Serial Number: CF424AE19N60085

Chassis Model: E234

Firmware version: 0.36

Location: Rack 2,1U

Power state: 

Drive List




Search for

Q

Asset Tag: Asset Tag

Erase On Detach: ☒ true

Modify Drive Configuration

ID	Slot ID	Capacity	Model	Manufac	Part Number	Serial Number	Asset Tag	Drive Era	Erased On t	Raw Dat
Zone: 14-ft-1-z-3 Composed Node ID: 2 System ID: 15-s-13 BMC IP: 192.168.5.2										
14-c-3-d-12	2	3726	P4500	INTEL	INTEL SSDPE2K	PHLF719500D74F			✓	
Drives available to attach										
14-c-3-d-10	0	3726	P4500	INTEL	INTEL SSDPE2K	PHLF723600XJ4F		✓	✓	
14-c-3-d-9	1	3726	P4500	INTEL	INTEL SSDPE2K	PHLF720401LB4F		✓	✓	

This page is intentionally left blank.


APPENDIX

SUPERMICRO BIGTWIN SERVERS

Supermicro's BigTwin is a high-performance, dense 2U compute platform with four independent compute nodes, two CPUs per node, powered by two redundant platinum-level power supplies.

For this deployment, each node has 192 GB of memory, two Intel M.2 SSDs, and 14 Intel 4 TB NVMe drives. While a standard BigTwin accommodates six NVMe drives available per node, the NVMe Enclosure increases this capacity. Data going into and coming out of each node use two 100 Gb Ethernet cards. Four nodes are assigned to the NVMe Enclosure, with each node accessing eight drives.

NVMe drives and Ethernet cards are balanced across both CPUs in the following way:



In essence: Each 100 Gb Ethernet card is attached to a separate CPU, NVMe Enclosure drives are connected to CPU1, and BigTwin Chassis NVMe's are connected to CPU2 for maximum use of available bandwidth.

NVME ENCLOSURE TECHNICAL DETAILS

Supermicro's 1U NVMe Enclosure is an intelligent device with 32 U.2 flash disks where up to 12 hosts may dynamically access any flash installed in the NVMe Enclosure.



Host to NVMe Enclosure connections require Mini-SAS HD connectors. Each cable provides a PCI-E x8 connection. For this deployment, two x8 cables are used to provide 16 PCI-E lanes per host. Four hosts are assigned to an NVMe Enclosure.

Figure 2 shows the numbering for the NVMe Enclosure host port connections.:




Figure 3. PCI-E Expansion Card

The connection towards the top of the card needs to go to the NVMe Enclosure's left connection; the connection towards the bottom of the card connects to the NVMe Enclosure's right connection. Pay attention to where the AOC is installed on the BigTwin. Figure 4 shows the physical connection of the PCI-E cables in the BigTwin Servers.




Figure 4. PCI-E Cabling to BigTwin Servers

After the host and NVMe Enclosure are connected, any attachment of drives to each host is made through the RSD POD Manager.

The NVMe Enclosure has four switches, two facing the upstream host connections and two facing the downstream drive connections. The four switches are grouped into a single logical switch by the enclosure firmware. Drives in each drive tray are attached to separate downstream switches. Systems with fewer than 32 drives should have the drives evenly distributed between the drive trays to maximize switch port usage. Drives should also be assigned across trays as well to maximize switch bandwidth utilization.

MOUNTING NVMe DRIVES

The NVMe architecture described in this document supports hot plug NVMe drives in a BigTwin chassis and an NVMe enclosure. The hot-plugging of NVMe drives will trigger OS PCI-E re-enumeration. This can result in device naming changes at the OS level. To avoid issues with mount points, volumes should be referenced by UID rather than device names in /etc/fstab.

Normally the format of /etc/fstab is

```
/dev/nvme0n1p1    /mnt/nvme0n1    ...
```

Using UUID requires executing `lsblk -f`, which will list the file system's UUID and /etc/fstab needs to reference it in the form.

```
UUID=<UUID from lsblk -f>    /mnt/nvme0n1 ...
```

PARTITIONING NVMe DRIVES

If drive locality awareness is required in the drive partitioning schema, the drive location can be identified by the PCI-E path (using `lspci`) or by the drive serial number.

- Using `lspci -t`:
Drives in the NVMe enclosure will be found in PCI-E paths that include intervening PCI-E switches.
- Using drive serial numbers:
The RSD Redfish API or POD Manager UI can be used to find the serial numbers of all drives in the enclosure. The drive collection can be retrieved to obtain the list of drives from the enclosure, followed by querying the details for each drive to obtain the serial numbers.
Drive Collection: `/redfish/v1/Chassis/151-c-2/Drives`
The serial number can be matched against the serial number reported at the OS.
`/sys/class/nvme/<drive>/serial`

DYNAMIC PCI-E ATTACHMENT/DETACHMENT

Linux has the option of permitting dynamic PCI device enumeration without OS restarts through the kernel parameter

```
pci=realloc
```

Normally these parameters are added in /etc/default/grub and `update-grub` is issued to commit the change. For automated Ubuntu deployments, the file can be edited using `sed` followed by `grub-mkconfig` to commit the change in a Kickstart configuration file:

```
%post
sed -ie 's/CMDLINE _LINUX _ DEFAULT="" /CMDLINE _LINUX _
```

```
DEFAULT="pci=realloc "/g' \ /etc/default/grub  
grub-mkconfig -o /boot/grub/grub.cfg  
%end
```

REFERENCE BOM

The following Bill of Materials represents the as-configured system used for validation of the design represented in this paper.

GROUP	PART	DESCRIPTION	USAGE
Rack/ Power	SRK-24SE-02	24U ENCLOSURE, 598x1000 (mm, WxD),HF	
	SRK-01PN-01	1U MESH BLANKING PANEL	
	SRK-01PN-11	1U SOLID BLANKING PANEL	
	MCP-210-00080-0B-01	1U Supermicro Logo Panel (for Rack)	
	SRK-01CM-41	1U BRACKET W/ FINGER C.M. TROUGH	Switch and JBoF Cable Mgmt
	SRK-01PD-76	APC AP9571A , Rack PDU, Basic, 1U, 30A, 208V, (10) C13	
	CBL-PWCD-0372-IS	PWCD,US,IEC60320 C14 TO C13,3FT,16AWG,RoHS/REACH	
	CBL-PWCD-0374-IS	PWCD,US,IEC60320 C14 TO C13,6FT,16AWG,RoHS/REACH	
	CBL-PWCD-0775	PWCD,US,NEMA5-15R TO IEC60320 C14,1FT,18AWG,HF,RoHS/REACH	
BigTwin	SYS-2029BT-HNR	X11DPT-B, CSE-217BHQ+-R2K22BP2, BPN-ADP-6NVME3-1UB	
	P4X-SKL6130-SR3B9	SKL 6130 4/2P 16C/32T 2.1G 22M 10.4GT 125W 3647 H0	
	MEM-DR416L-CL06-ER26	16GB DDR4-2666 2Rx8 on-die ECC RDIMM	
	AOC-SLG3-2M2-P	Low Profile PCI-E Riser Card supports 2 M.2 Module, RoHS	Local boot
	HDS-M2M-SSDPEKKA128G7	Intel DC P3100 128G NVMePCI-E3.0M.2 22x80mm0.3DWPDP,HF,RoH	
	AOC-MH25G-m2S2T	SIOM Dual-port 25GbE controller with 2 SFP28 ports based on Mellanox ConnectX-4 Lx EN; and Dual-port 10GbE RJ45 (10GBase-T) based on Intel X550-AT2	
	AOC-SLG3-4X4P-P	PCI-Ex16, Low profile 4x external port NVMe PLX switch, HF, RoHS	
	SFT-DCMS-SINGLE	Supermicro System Management Software Suite node license, HF, RoHS/REACH, PBF	

GROUP	PART	DESCRIPTION	USAGE
Mgmt/ Network	SYS-5019S-TN4-SRSMGT and SFT-SMCI-SRSDM	1U X11 SUPERSERVER Rack Management Server HW and RSD SW	PODM/RMM
	SFT-DCMS-SINGLE	Supermicro System Management Management Software Suite node license, HF, RoHS/REACH, PBF	
	SSE-G48-TG4	1/10G STANDALONE SWITCH-48 1G PORT 4 10G PORTS	MgmtMgmt.
	SSE-X3348TR	48Port 10GBASE-T Switch-1U Rack W/4-Port 40GbE Nor. Air	Data
	SSE-C3632SR	32-port 100GbE QSFP28,B2F airflow,2x800W	NVMe Fabric
	SFT-CLSPL100G-1Y	Cumulus-Linux SW 100G Perpetual License with 1 yyr. Cumulus SnS	
	CSE-PT52	THIRD GENERATION GUIDE RAIL ASSY. Outer slides extendable length: 26" to 33.5"	Switch rails
	CBL-C6-GN3FT-J	ETHERNET, CAT6, RJ45, SNAGLESS, GREEN, UTP, 3FT(0.9M), 28AWG, RoHS	Appliance MgmtMgmt.
	CBL-0362L	ETHERNET, CAT6, RJ45 W/BOOT, YELLOW, UTP, 2FT (60CM), 24AWG	MgmtMgmt./ Data link
	CBL-NTWK-0988-QS28C20M-1	ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28,2m,30 AWG, Mellanox	100Gb Data
	CBL-0364L	ETHERNET, CAT6, RJ45 W/BOOT, YELLOW, UTP, 4FT (1.2M), 24AWG	LAN1
	CBL-0365L	ETHERNET, CAT6, RJ45 W/BOOT, YELLOW, UTP, 5FT (1.5M), 24AWG	LAN1
	CBL-NTWK-0598	ETHERNET, CAT6, RJ45, SNAGLESS, BLUE, UTP, 4FT (1.2M),24AWG	LAN2
	CBL-NTWK-0535	ETHERNET, CAT6, RJ45, W/BOOT, BLUE, UTP, 5FT (1.5M), 24AWG	LAN2
	CBL-C6-GN4FT-J	ETHERNET, CAT6, RJ45, SNAGLESS, GREEN, UTP, 4FT(1.2M), 28AWG	BMC
	CBL-C6-GN5FT-J	ETHERNET, CAT6, RJ45, SNAGLESS, GREEN, UTP, 5FT(1.5M), 28AWG	BMC
NVMe Enclosure	SSG-136R-N32JBF	1U 32 NVMe JBOF	
	HDS-2VD-SSDPE2KX040T701	Intel DC P4500 4TB NVMe PCI-E 3.0 3D TLC 2.5" 1DWPD FW130	
	CBL-SAST-1035-1	MINI SAS HD, x8, 12G,EXT, PASSIVE, PULL, 1M, 30AWG, RoHS	

About Super Micro Computer, Inc.

Supermicro® (NASDAQ: SMCI), the leading innovator in high-performance, high-efficiency server technology is a premier provider of advanced server Building Block Solutions® for Data Center, Cloud Computing, Enterprise IT, Hadoop/Big Data, HPC and Embedded Systems worldwide. Supermicro is committed to protecting the environment through its “We Keep IT Green™” initiative and provides customers with the most energy-efficient, environmentally-friendly solutions available on the market.

Learn more on www.supermicro.com

No part of this document covered by copyright may be reproduced in any form or by any means — graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system — without prior written permission of the copyright owner.

Supermicro, the Supermicro logo, Building Block Solutions, We Keep IT Green, SuperServer, Twin, BigTwin, TwinPro, TwinPro², SuperDoctor are trademarks and/or registered trademarks of Super Micro Computer, Inc.

Ultrabook, Celeron, Celeron Inside, Core Inside, Intel, Intel Logo, Intel Atom, Intel Atom Inside, Intel Core, Intel Inside, Intel Inside Logo, Intel vPro, Itanium, Itanium Inside, Pentium, Pentium Inside, vPro Inside, Xeon, Xeon Phi, and Xeon Inside are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.

All other brands names and trademarks are the property of their respective owners.

© Copyright 2018 Super Micro Computer, Inc. All rights reserved.

