

# SUPERMICRO®

# Apache™ Hadoop®

*High Performance Enterprise Hadoop Infrastructure*




1U SC815 - Name Nodes



SC826 - 12 Drives



4U Fat Twin™ FT4- 8 Drives per Node



SC825 - 8 Drives



2U Twin SC827 - 6 Drives per Node



10GbE/1GbE Switches

## The industry's broadest line of Enterprise Hadoop Infrastructure & Server Building Block Solutions®

- Wide range of servers based upon Intel® Xeon® Processor E5-2600 series
- Enterprise-ready compute and storage with high reliability, quality and performance at an affordable price point
- Hot Swap Direct Attach 4x, 6x, 8x, and 12x 3.5" Hard Disk Drives per Node
- Up to Platinum Level (95%+) Efficiency - Redundant Power Supplies with PMBus
- Remote Systems and Cluster Management with Onboard IPMI and included IPMI View
- Cost effective Layer 3 10GbE Switches and Layer 3 1GbE Switches with 10GbE uplinks

[www.supermicro.com/hadoop](http://www.supermicro.com/hadoop)

March 2013



MODEL	815 Name Node /Job Tracker	827 2U Twin Data Node (6x 3.5" HDDs per Node)	825 Data Node (8x 3.5" HDDs)
<b>Key Features</b>	<ul style="list-style-type: none"> <li>Optimized Hadoop Name Node</li> <li>IPMI 2.0 + KVM with dedicated LAN for remote management</li> <li>Cost effective, high performance, power efficient</li> <li>Flexible WIO controller or I/O add-on cards</li> <li>4x3.5" hot-swap drive bays</li> <li>Redundant 700W <b>Gold Level</b> (93%) high-efficiency power supplies</li> </ul>	<ul style="list-style-type: none"> <li>Optimized Hadoop Twin Data Node</li> <li>IPMI 2.0 + KVM with dedicated LAN for remote management</li> <li>Cost effective, high performance, power efficient</li> <li>Two 6x HDD Nodes in 2U of rack space</li> <li>6x 3.5" hot-swap drive bays per node</li> <li>Redundant 920W <b>Platinum Level</b> (94%) high-efficiency power supplies</li> </ul>	<ul style="list-style-type: none"> <li>Optimized Hadoop Data Node</li> <li>IPMI 2.0 + KVM with dedicated LAN for remote management</li> <li>Cost effective, high performance, power efficient</li> <li>8x 3.5" hot-swap drive bays</li> <li>Redundant 740W <b>Platinum Level</b> (94%) high-efficiency power supplies</li> </ul>
<b>Processor Support</b>	Intel® Xeon® processor E5-2600	Intel® Xeon® processor E5-2600	Intel® Xeon® processor E5-2600
<b>Chipset</b>	Intel® 602 chipset with QPI up to 8.0GT/s	Intel® 602 chipset with QPI up to 8.0GT/s	Intel® 602J chipset with QPI up to 8.0GT/s
<b>System Memory (Max)</b>	<b>64GB</b> (Max. 512 GB in 16 x DDR3 Reg. ECC RDIMM)	<b>32GB</b> (Max. 256 GB) of DDR3 Reg. ECC Per Node (2 nodes per system)	<b>32GB</b> (Max. 512 GB in 16 x DDR3 Reg. ECC RDIMM)
<b>Expansion Slots</b>	2x PCI-E 3.0 x16	3x PCI-E 3.0 x 8 slots	6x PCI-E 3.0 x 8
<b>OnBoard &amp; AOC HBA Controllers</b>	Intel® 602 SCU SATA RAID 0, 1, 5, 10 support AOC - Hardware RAID LSI 2208 controller RAID 0, 1, 5, 6, 10	Intel® 602 SCU and AHCI SATA RAID 0, 1, 5, 10 support	Intel® 602J SCU SATA RAID 0, 1, 5, 10 support LSI 2308 controller IT Mode
<b>Networking /VGA/Audio</b>	Intel® i350 Dual Port Gigabit Ethernet Virtual Machine Device Queues reduce I/O overhead G200 (Renesas SH7757 BMC)	Intel® i350 Dual Port Gigabit Ethernet Virtual Machine Device Queues reduce I/O overhead G200 (Renesas SH7757 BMC)	Intel® i350 Quad/4 Port Gigabit Ethernet Virtual Machine Device Queues reduce I/O overhead G200 (Renesas SH7757 BMC)
<b>Manageability</b>	Dedicated IPMI 2.0 with virtual media over LAN and KVM-over-LAN support	Dedicated IPMI 2.0 with virtual media over LAN and KVM-over-LAN support	Dedicated IPMI 2.0 with virtual media over LAN and KVM-over-LAN support
<b>Drive Bays</b>	4x 3.5" SAS/SATA hot-swap HDD bays - 4x 300GB <b>15K</b> HDD	12x 3.5" SATA hot-swap HDD bays 6x 3.5" HDD per node	8x 3.5" SAS/SATA hot-swap HDD bays
<b>Power Supply</b>	Redundant (1+1) 700W <b>Gold Level</b> (93%) high-efficiency power supplies w/ PMBus	Redundant (1+1) 920W <b>Platinum Level</b> (94%) high-efficiency power supplies w/ PMBus	Redundant (1+1) 740W <b>Platinum Level</b> (94%) high-efficiency power supplies w/ PMBus
<b>Cooling System</b>	4x 4cm heavy duty counter-rotating fans	4x 8cm heavy duty fans with PWM fan speed control	3x 8cm 6300rpm hot-swap PWM cooling fans
<b>Form Factor</b>	1U rackmount chassis (25.6" depth) 17.2"W x 1.7"H x 25.6"D (437 x 43 x 650mm)	2U rackmount chassis (28.5" depth) 17.25"W x 3.47"H x 28.5"D (438 x 88 x 724mm)	2U rackmount chassis (25.5" depth) 17.2"W x 3.5"H x 25.5"D (437 x 89 x 648mm)





cloudera  
**CONNECT**  
CERTIFIED TECHNOLOGY



cloudera  
**CONNECT**  
CERTIFIED TECHNOLOGY



## FT4 Data Node (8x 3.5" HDDs per Node)

- Optimized Hadoop Data Node
- IPMI 2.0 + KVM with dedicated LAN for remote management
- Cost effective, high performance, power efficient
- 4 Nodes of 8 x HDD per node in 4U of rack space
- 32 x 3.5" hot-swap drive bays
- Redundant 4x1280W **Platinum Level** (95%+) high-efficiency power supplies

Intel® Xeon® processor E5-2600

Intel® 602 chipset with QPI up to 8.0GT/s

**32GB** (Max. 512 GB) of DDR3 Reg. ECC per Node

1x PCI-E 3.0 x 8  
1x MicroLP PCI-E 3.0 x 8

Intel® 602 SCU and AHCI  
SATA RAID 0, 1, 5, 10 support

Intel® i350 Dual Port Gigabit Ethernet  
Virtual Machine Device Queues reduce I/O overhead  
G200 (Renesas SH7757 BMC)

Dedicated IPMI 2.0 with virtual media  
over LAN and KVM-over-LAN support

32x 3.5" SATA hot-swap HDD bays  
8x 3.5" HDD per Node

Redundant 4x1280W **Platinum Level** (95%+) high efficiency power supplies w/ I2C & PMBus

3x 8cm 6300rpm hot-swap PWM cooling fans

4U rackmount chassis (29" depth)  
17.63"W x 6.96"H x 29"D (448 x 177 x 737mm)

## 826 Data Node (12x 3.5" HDDs)

- Optimized Hadoop Data Node
- IPMI 2.0 + KVM with dedicated LAN for remote management
- Cost effective, high performance, power efficient
- 12x 3.5" hot-swap drive bays
- Redundant 740W **Platinum Level** (94%) high-efficiency power supplies

Intel® Xeon® processor E5-2600

Intel® 602 chipset with QPI up to 8.0GT/s

**32GB** (Max. 512 GB) of DDR3 Reg. ECC

6x PCI-E 3.0 x 8

Intel® 602J SCU and AHCI  
SATA RAID 0, 1, 5, 10 support  
LSI 2308 controller IT Mode

Intel® i350 Quad/4 Port 1G Gigabit Ethernet  
Virtual Machine Device Queues reduce I/O overhead  
G200 (Renesas SH7757 BMC)

Dedicated IPMI 2.0 with virtual media  
over LAN and KVM-over-LAN support

12x 3.5" SAS/SATA hot-swap HDD bays  
2x 2.5" SAS/SATA HDD bays

Redundant (1+1) 740W **Platinum Level** (94%) high-efficiency power supplies w/ PMBus

3x 8cm hot-swap cooling PWM fans

2U rackmount chassis (25.5" depth)  
17.2"W x 3.5"H x 25.5"D (437 x 89 x 648mm)

## 1GbE & 10GbE Ethernet Switches

### Key Features

- Layer 2 / 3 Ethernet Switching
- Link aggregation (LACP)
- Jumbo frames
- VLAN support

### Configuration Flexibility

#### • SSE-G24-TG4

- 24 ports of RJ45 1Gbps Ethernet with 4 SFP Combo ports
- 4 ports 10Gbps Ethernet (CX4, XFP, or SFP+)
- Stacking

#### • SSE-G48-TG4

- 48 ports of RJ45 1Gbps Ethernet with 4 SFP Combo ports
- 4 ports 10Gbps Ethernet (CX4, XFP, or SFP+)
- Stacking

#### • SSE-X24S / SSE-X3348S(R)

- 24 / 48 ports 10Gbps Ethernet (SFP+)
- Reverse airflow option

### Management

- Web-based management interface
- Industry standard CLI with telnet, SSH or local management port
- Scripting
- Logging

### Routing

#### Multicast

- Quality of service and DiffServ
  - 8 priority queues per port
  - Adjusted WRR and Strict Priority Scheduling
- Security
  - 802.1x port based authentication
  - RADIUS and TACACS+ authentication
  - SSH / SSL encryption

### Dimensions


- 17.3"W x 15.2"D x 1.7"H (440 x 386 x 43mm)

For More Information:

<http://www.supermicro.com/products/nfo/networking.cfm#switch>

### IPMI Server Management

- HW fault alerts – E-mail or SNMP
- Redundant access for management
- Support for open standards (SMASH 2.0, LDAP, WSMAN)
- VLAN support
- KVM and Virtual Media and SOL (no additional license needed)
- Easy maintenance (Command line flash tools for firmware update)




For More Information:

<http://www.supermicro.com/products/nfo/IPMI.cfm>

**H**adoop is an open-source project administered by the Apache Software Foundation. Hadoop's contributors work for some of the world's biggest technology companies. That diverse, motivated community has produced a genuinely innovative platform for consolidating, combining and understanding data. Enterprises today collect and generate more data than ever before. Relational and data warehouse products excel at OLAP and OLTP workloads over structured data. Hadoop, however, was designed to solve a different problem: the scalable, reliable storage and analysis of both structured and complex data. As a result, many enterprises deploy Hadoop alongside their legacy IT systems, allowing them to combine old and new data sets in powerful new ways. Technically, Hadoop consists of two key services: reliable data storage using the Hadoop Distributed File System (HDFS) and high-performance parallel data processing using a technique called MapReduce. Hadoop runs on a collection of commodity, shared-nothing servers. You can add or remove servers in a Hadoop cluster at will; the system detects and compensates for hardware or system problems on any server. Hadoop, in other words, is self-healing. It can deliver data - and can run large-scale, high-performance processing jobs - in spite of system changes or failures.

## Why Are Supermicro's Turnkey Cluster Racks ideal for Apache<sup>™</sup> Hadoop<sup>®</sup>?

Supermicro designs and develops the ideal turnkey pilot racks for getting started with Apache<sup>™</sup> Hadoop<sup>®</sup>. Leveraging Supermicro's optimized servers and switches as a foundation, Supermicro has designed two turnkey racks to get anyone started - 14U and 42U versions. Supermicro has focused on integration, remote systems and power management to lower the deployment and commissioning timeframes to provide developers access to native hardware environments, quickly and easily.



14U Rack



42U Rack

### 14U Hadoop Rack

- Optimized 14U Hadoop Rack
- IPMI 2.0 + KVM with dedicated LAN for remote management
- Cost effective, high performance, power efficient server nodes
- 2 Primary / Secondary Name Nodes
- 64 GB Memory Per Node
- 8 Data Nodes
- 32 GB Memory Per Node
- 48TB Raw HDFS Capacity
- 2x Intel<sup>®</sup> Xeon<sup>®</sup> E5-2600 Processors per Node
- Platinum Level (94%) efficient power supplies

### 14U Hadoop Rack BOM

- 2x SYS-HNW0-15026364-HADP
- 4x SYS-HDT0-27126332-HADP
- 1x SSE-G48-TG4 (48 Ports)
- 1x SC512 1U Server (Cluster SSM Server Management)
- 1x 2U PDU
- Integrated, Cabled, Labeled, Tested, Crated and Shipped as turnkey rack PN# SRS-14UL63-HADP-TL

### 42U Hadoop Rack

- Optimized 42U Hadoop Rack
- IPMI 2.0 + KVM with dedicated LAN for remote management
- Cost effective, high performance, power efficient server nodes
- 3x Primary / Secondary Name Nodes
- 64 GB Memory Per Node
- 36x Data Nodes
- 32 GB Memory Per Node
- 216-1728TB Raw HDFS Capacity
- 2x Intel<sup>®</sup> Xeon<sup>®</sup> E5-2600 Processors per Node
- Platinum Level (94-95%) efficient power supplies

### 42U Hadoop Rack BOM

- 3x SYS-HNWX-19026364-HADP
- 18x SYS-HDTX-27126332-HADP
- 1x SSE-G48-TG4 (48 1G Ports)
- 1x SSE-X3348SR (48 Ports 10G SFP+)
- 1x SC512 1U Server (Cluster SSM Server Management)
- 4x PDU
- Integrated, Cabled, Labeled, Tested, Crated and Shipped as turnkey rack PN# SRS-42XL63-HADP-TL
- 42U Fat Twin Hadoop Lead Rack PN# SRS-42XL63-HADP-FL

## Optimized Hadoop Nodes



Supermicro's extensive line of products include 1U, 2U and 4U rack mount servers optimized for Hadoop. Customers can order the following configurations from their local Supermicro reseller. Supermicro Hadoop systems feature the industry's leading thermal and power designs to deliver optimal performance-per-watt and performance-per-dollar.

Supermicro Model #	Description
SYS-HNW0-15026364-HADP	1U Hadoop 815 Name Node 6C E5-2630 2.3G 15M 7.2G 95W 64GB MEM
SYS-HDT0-27126332-HADP	2U Hadoop 827 <b>6</b> 1TB 6C E5-2630 2.3G 15M 7.2GT 95W 32GB MEM
SYS-HDT0-27226332-HADP	2U Hadoop 827 <b>6</b> 2TB 6C E5-2630 2.3G 15M 7.2GT 95W 32GB MEM
SYS-HDF0-24126332-HADP	4U Hadoop FT4 <b>8</b> 1TB 6C E5-2630 2.3G 15M 7.2GT 95W 32GB MEM
SYS-HDF0-24126332-HADP	4U Hadoop FT4 <b>8</b> 2TB 6C E5-2630 2.3G 15M 7.2GT 95W 32GB MEM
SYS-HDD0-25126332-HADP	2U Hadoop 825 <b>8</b> 1TB 6C E5-2630 2.3G 15M 7.2GT 95W 32GB MEM
SYS-HDD0-25226332-HADP	2U Hadoop 825 <b>8</b> 2TB 6C E5-2630 2.3G 15M 7.2GT 95W 32GB MEM
SYS-HDD0-26126332-HADP	2U Hadoop 826 <b>12</b> 1TB 6C E5-2630 2.3G 15M 7.2GT 95W 32GB MEM
SYS-HDD0-26226332-HADP	2U Hadoop 826 <b>12</b> 2TB 6C E5-2630 2.3G 15M 7.2GT 95W 32GB MEM
SYS-HNWX-19026364-HADP	1U Hadoop 819 Name Node 6C E5-2630 2.3G 64GB MEM 10GSFP+
SYS-HDTX-27126332-HADP	2U Hadoop 827 <b>6</b> 1TB 6C E5-2630 2.3G 32GB MEM 10GSFP+
SYS-HDFX-24126332-HADP	4U Hadoop FT4 <b>8</b> 1TB 6C E5-2630 2.3G 32GB MEM 10GSFP+
SYS-HDDX-26126332-HADP	2U Hadoop 826 <b>12</b> 1TB 6C E5-2630 2.3G 32GB MEM 10GSFP+

[www.supermicro.com/hadoop](http://www.supermicro.com/hadoop)