Supermicro: Worldwide Leader in GP/GPU Servers and Workstation Platforms

GTC May 16, 2012
Presented by: Don Clegg
VP, Marketing & Business Development

SUPERMICRO®

From GTC Conference Session Abstract:

Discover the measurable advantages
that make Supermicro the
Time-to-Market Leader
in GPU platform enablement.

See how Supermicro's innovative
Application-Optimized designs
enable partners to both
scale-up and scale-out
for maximum return on investment.

Review platforms that highlight
Supermicro's leadership in
Compute Density,
Peak Performance,
Scalability,
Power Efficiency,
Manageability,
Reliability and
Cost Effectiveness

What's Presented

- **Corporate Overview Update**
  - Time-to-Market Leadership
- **Application - Optimized**
  - Scale-Up / Scale-Out
- **Supermicro Product Leadership**
- **Customer Endorsements**

Key Questions to Answer

➔ **Who are these guys… really?**

- What's the secret to their Time-to-Market Leadership?
- How do they consistently out-innovate other Tier-Ones?
- Why will the trend continue for the foreseeable future?
- Can they really meet my rigorous product and logistic requirements?

➔ **Why should I care if my design is Application-Optimized?**

- Is there ONE ideal GP / GPU platform for all applications?
- Which GPU platforms are best for Scaling-Up / -Out?

➔ **Are the product benefits measurable?**

- Which products are GPU optimized?
- What are specific examples of product leadership?

➔ **Who are some of Supermicro's partners?**

- Why have they chosen Supermicro?

- **Corporate Overview Update**
  - Time-to-Market Leadership

➔ **Who are these guys… really?**

- What's the secret to their Time-to-Market Leadership?
- How do they consistently out-innovate other Tier-Ones?
- Why will the trend continue for the foreseeable future?
- Can they really meet my rigorous product and logistic requirements?

Corporate Overview / Update

Why Supermicro has the Time-to-Market Advantage

Supermicro Overview

- Founded in 1993, headquartered in Silicon Valley, USA (NASDA 2007: SMCI)
- Time-to-Market leader in server technology innovation and green computing
- Broadest server portfolio in the industry;
  Designs are In-House / "under the same roof"
- Supermicro's application-optimized, high-efficiency servers, GPU systems, networks, storage and workstations are deployed globally across data centers, HPC and critical IT infrastructures

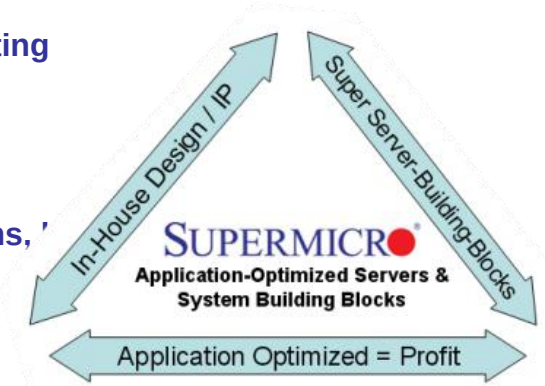

The Supermicro Time-to-Market Advantage

Server Building Block Solutions ®

Resource Optimized

Customer Driven Design Expertise

Application Optimized

SUPERMICRO®

Server Building Block Solutions®

>550 Motherboards (1)	>1300 Chassis (1)	> 350 Cooling Modules (1)	> 140 Power Supplies (1)	Open CPU/ Memory	Operating Systems / Applications
SUPERMICRO	SUPERMICRO	SUPERMICRO	SUPERMICRO	Industry Standard (AMD, intel, NVIDIA)	Customer Choice


Blade Servers


Twin Architecture


GPU Solutions


Fat Twin


Storage


Embedded


Switches

Positioned for Continued Growth

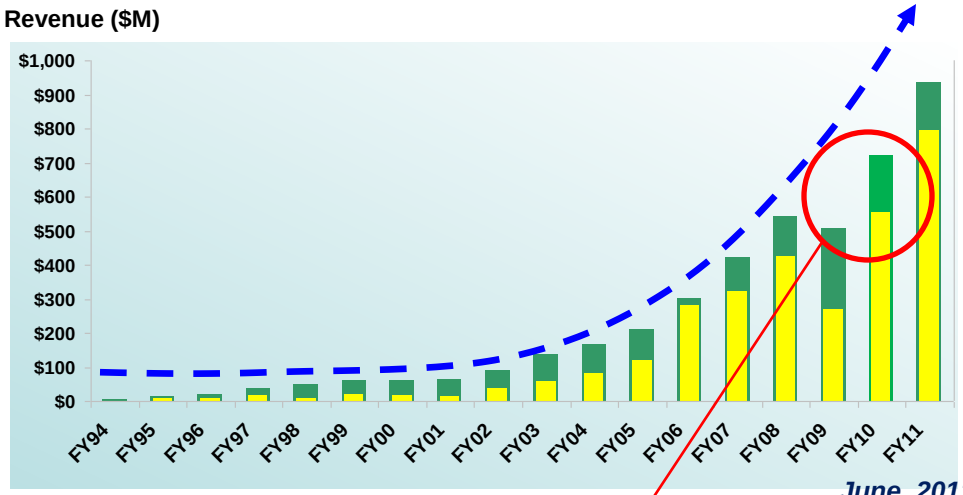

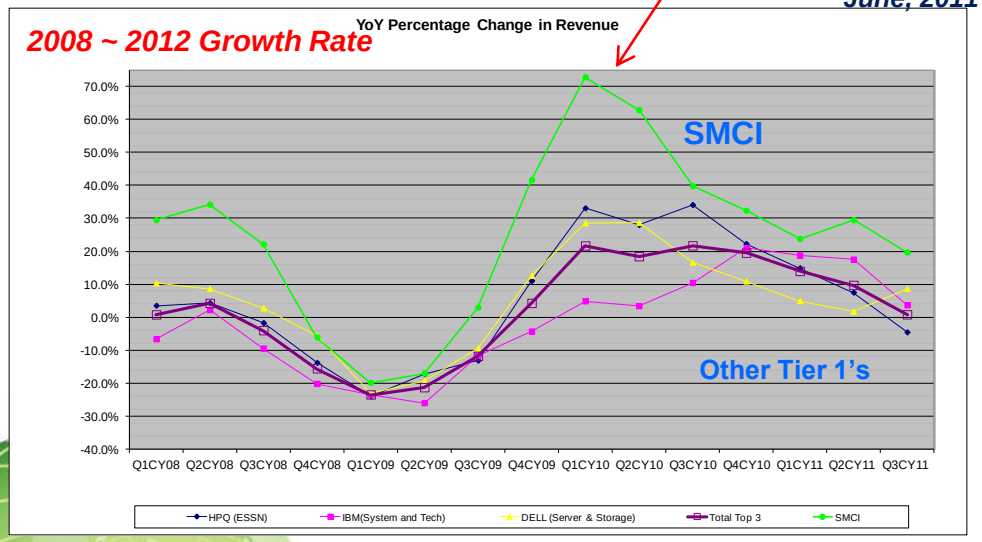

- Profitable every year since founding, no exception
- Consistently outgrowing our aggressive competitors
- Last to suffer from 2009 economic tsunami; First and Fastest to rebound (with the highest percentage recovery)
- Winning with Superior Technology and Consistent Time-to-Market Leadership
- 2012: Positioned to grow even faster through increased scale

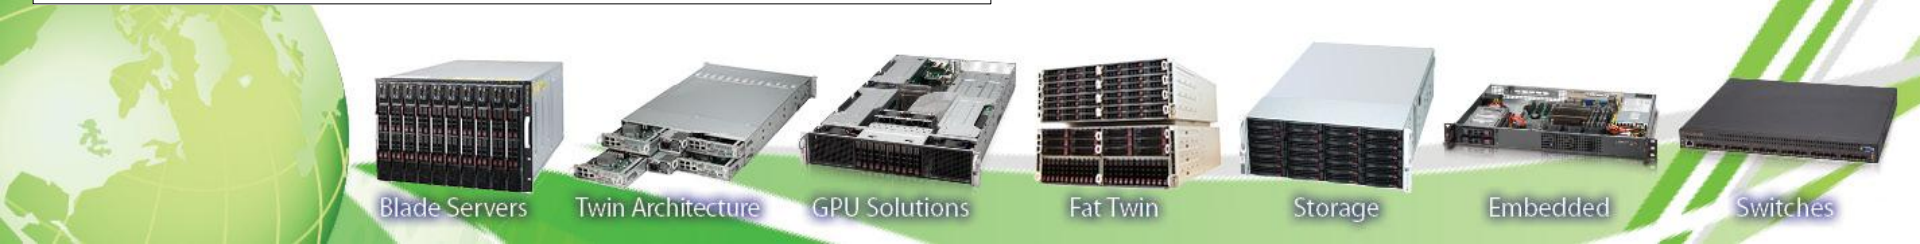

Supermicro Global Support

Main Campus
(San Jose)

Fremont Facility

Worldwide Headquarters | European Operations Center | Asian Technology Park

- Corporate locations (US, Taiwan, Netherlands)
- Local supports in more than 70 countries
- Distribution and VAR alliances
- Commitment to help customers get the most optimized HPC, Cloud and Datacenter Solutions

Supermicro Science & Technology Park

1.6 Million Square Feet

Comprehensive Manufacturing Functionality
R&D, OEM/ODM Partnership, System/Solution Validation, Production/Assembly, Logistics, RMA

30K Server System/ month new integration capacity
2x previous manufacturing capacity

Recap

- **Corporate Overview Update**
  - Time-to-Market Leadership

➔ **Who are these guys… really?**

- What's the secret to their Time-to-Market Leadership?
- How do they consistently out-innovate other Tier-Ones?
- Why will the trend continue for the foreseeable future?
- Can they really meet my rigorous product and logistic requirements?

Section

Q & A

- **Application - Optimized**
  - Scale-Up / Scale-Out

➔ **Why should I care if my design is Application-Optimized?**

- Is there ONE ideal GP / GPU platform for all applications?
- Which GPU platforms are best for Scaling-Up / -Out?

Application Optimized

Scale-Up / Scale-Out Product Overview

Application Optimization

Server	Application	Optimization
MP	Virtualization	Mission Critical
Twin	SMB	Scalability
Blade	Datacenter	Efficiency
GPU	HPC	Computation
Embedded	Single Purpose	Reliability
Storage	NAS, SAN	Capacity
WIO/UIO	Networking	Flexibility

Ideal Solution

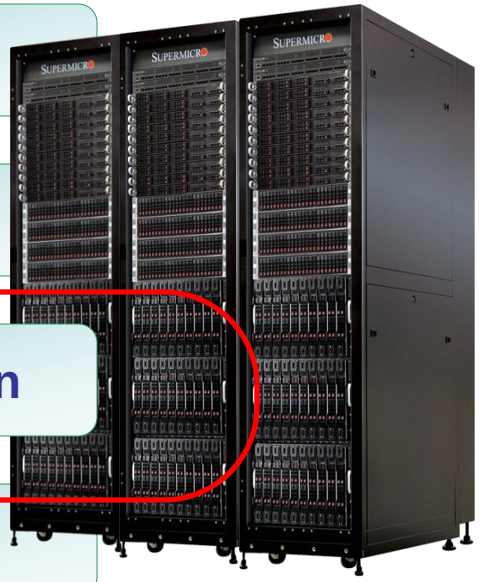

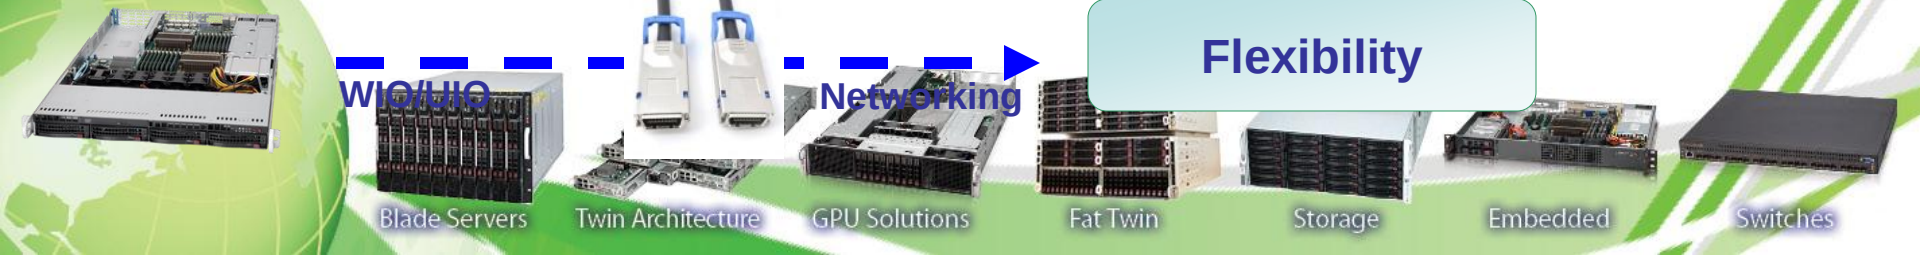

MOST EXTENSIVE X9 (E5 SANDYBRIDGE) GPU SOLUTIONS AVAILABLE NOW

- **Best Products / Global Logistics / Time-To-Market**
- **Application Optimized Building Block Solutions**

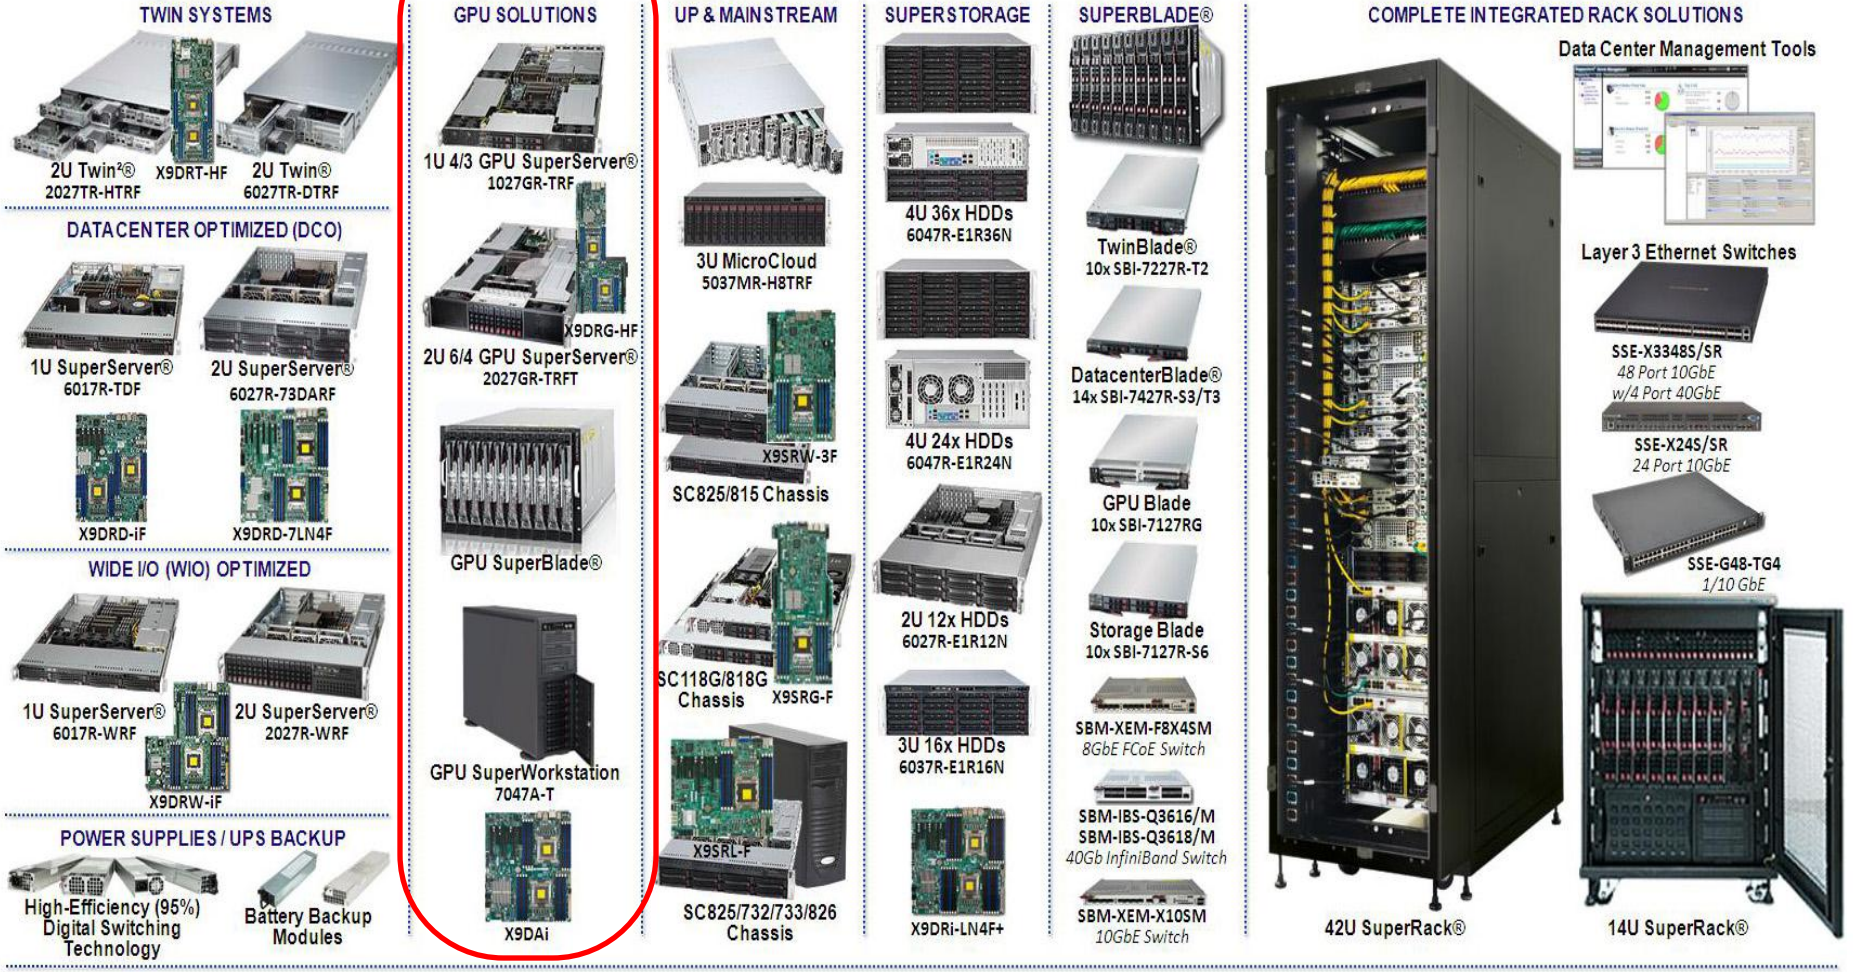

Leading Application Optimized Solutions

- **HPC/Enterprise 8-way/4-way Systems**
  - 8-way 5U high-end high-margin solution – 80 cores and 2TB
  - Nehalem-EX MP and G34 MP system in 1U/2U/4U
- **Storage Product Lines**
  - Double-Sided™, Super SBB product line with software solutions
- **SuperRack™**
  - Double side access, cable management, water cooling
  - Hadoop solution
- **Switch Products: 10GbE, IB, FCoE**
  - 10GbE onboard and 10GbE standalone switch
  - FCoE solution coming soon
- **MicroCloud™**
  - Optimized for IaaS
- **Power Subsystems**
  - High efficiency (94%+), digital switching, UPS w/ battery
- **Software Solutions**
  - Remote management, power management (NMview & SSM)
  - HPC/DC management toolset
  - Window OS Integration/Bundle

5U / 8-way

1U / 4-way

Double-Sided Storage

Front

Back

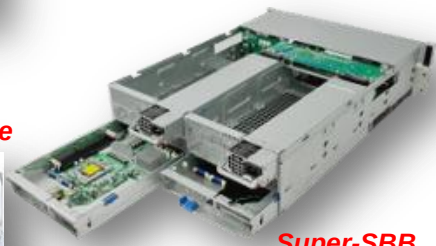

Super-SBB

SuperRack

MicroCloud

High Efficiency Power

Blade Servers

Twin Architecture

GPU Solutions

Fat Twin

Storage

Embedded

Switches

Leading Solutions - Snapshot

- **GPU optimized product lines (1U, 2U and 4U)**
  - 1U, 2U New High Density GPU product line
- **SuperBlade™ - GPU Blade and TwinBlade™**
  - GPU blade (20GPUs in 7U, most dense in the industry)
  - 20/28 DP nodes in 7U, support 40Gb/s Infiniband or 10G Ethernet connectivity
- **X9 Sandy Bridge Solutions**
- **H8 Interlagos Solutions**
- **Twin Architecture**
  - 2U Twin (6x 3.5" HDD or 12x 2.5" HDD per node)
  - 1U Twin, 2U Twin², 2U Twin³ (8 nodes in 2U)
  - Fat Twin Architecture coming soon
- **IPC and Embedded Applications**
  - Atom and Core-based low power server: fan-less / long life cycle, for embedded and server appliances
- **Workstation and High-end Desktop Solutions**
  - Sandy Bridge UP, Near-silent in operation (21dB)
  - Everest Solution : High frequency trading (HFT) application

Twin Architectures

GPU Blade

TwinBlade™

GPU-optimized

Industrial PC

Workstation

Technology Progression – Application Optimization

Telsa S1070

1U 4-GPU Standalone box

PCI-E x16 PCI-E x16

6016TT-TF

1U Twin™

The fastest 1U server in the world

The first optimized CPU/GPU integrated Hybrid System

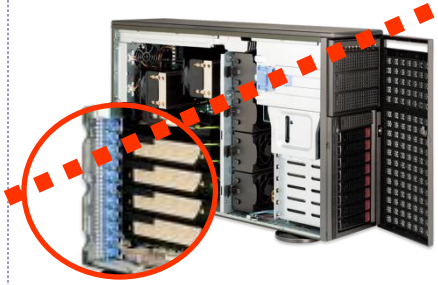

SW7046GT-TRF

The most powerful PSC

Next generation CPU & GPU support

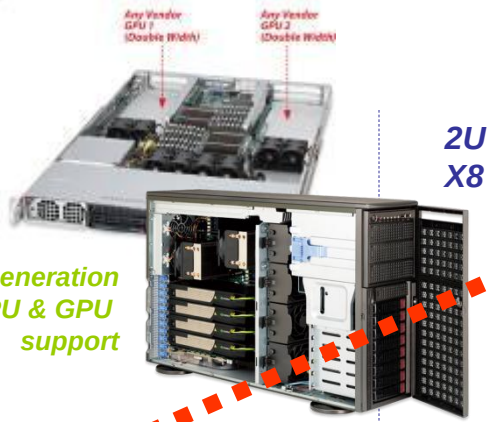

2U GPU / with QDR IB onboard

2U Twin

SC827HD-R1400B

front view

2U 6-GPU Integrated X8 and X9

1U 4-GPU Integrated X8 and X9

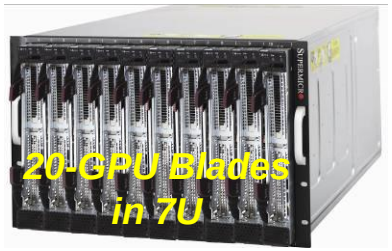

Working with the right technology partner is the key!

2008 2009 2010/11 2012/13


Blade Servers


Twin Architecture


GPU Solutions


Fat Twin


Storage


Embedded


Switches

Recap

- **Application - Optimized**
  - Scale-Up / Scale-Out

➔ **Why should I care if my design is Application-Optimized?**

- Is there ONE ideal GP / GPU platform for all applications?
- Which GPU platforms are best for Scaling-Up / -Out?

Section

Q & A

- **Supermicro Product Leadership**

➔ **Are the product benefits measurable?**

- Which specific products are GPU optimized?
- What are specific examples of product leadership?

Supermicro Product Leadership

MOST EXTENSIVE X9 (E5 SANDYBRIDGE) GPU SOLUTIONS AVAILABLE NOW

- **Best Products / Global Logistics / Time-To-Market**
- **Application Optimized Building Block Solutions**

TWIN SYSTEMS

2U Twin²® X9DRT-HF
2027TR-HTRF
2U Twin®
6027TR-DTRF

DATACENTER OPTIMIZED (DCO)

1U SuperServer®
6017R-TDF
2U SuperServer®
6027R-73DARF

X9DRD-iF
X9DRD-7LN4F

WIDE I/O (WIO) OPTIMIZED

1U SuperServer®
6017R-WRF
2U SuperServer®
2027R-WRF
X9DRW-iF

POWER SUPPLIES / UPS BACKUP

High-Efficiency (95%) Digital Switching Technology
Battery Backup Modules

GPU SOLUTIONS

1U 4/3 GPU SuperServer®
1027GR-TRF

X9DRG-HF
2U 6/4 GPU SuperServer®
2027GR-TRFT

GPU SuperBlade®

GPU SuperWorkstation
7047A-T

X9DAi

UP & MAINSTREAM

3U MicroCloud
5037MR-H8TRF

X9SRW-3F
SC825/815 Chassis

SC118G/818G Chassis
X9SRG-F

X9SRL-F
SC825/732/733/826 Chassis

SUPERSTORAGE

4U 36x HDDs
6047R-E1R36N

4U 24x HDDs
6047R-E1R24N

2U 12x HDDs
6027R-E1R12N

3U 16x HDDs
6037R-E1R16N

X9DRi-LN4F+

SUPERBLADE®

TwinBlade®
10x SBI-7227R-T2

DatacenterBlade®
14x SBI-7427R-S3/T3

GPU Blade
10x SBI-7127RG

Storage Blade
10x SBI-7127R-S6

SBM-XEM-F8X4SM
8GbE FCoE Switch

SBM-IBS-Q3616/M
SBM-IBS-Q3618/M
40Gb InfiniBand Switch

SBM-XEM-X10SM
10GbE Switch

COMPLETE INTEGRATED RACK SOLUTIONS

42U SuperRack®

Data Center Management Tools

Layer 3 Ethernet Switches

SSE-X3348S/SR
48 Port 10GbE w/4 Port 40GbE

SSE-X24S/SR
24 Port 10GbE

SSE-G48-TG4
1/10 GbE

14U SuperRack®

X9 (E5-SandyBridge) GPU LINEUP AVAILABLE NOW

GPU Servers

SuperServer
SYS-2027GR-TRF
SYS-2027GR-TRF-FM407
SYS-2027GR-TRF-FM409

SuperServer
SYS-1027GR-TRF
SYS-1027GR-TRF-FM307
SYS-1027GR-TRF-FM309

SuperServer
SYS-5017GR-TF

SuperServer
SYS-1017GR-TF

GPU Workstations

SuperWorkstation
SYS-7047GR-TRF
SYS-7047GR-TPRF
SYS-7047GR-TRF-FC409
SYS-7047GR-TPRF-FM407
SYS-7047GR-TPRF-FM409

GPU Blades

SuperBlade
SBI-7127RG + SBE-720E

* Pictures are reference only

Pre-configured Turnkey GPU Simclusters

** Results based on Tesla benchmarks. Kepler performance will be higher.*

	SRS-14URKS-GPUS-11	RS-14URKS-GPUS-12	SRS-42URKS-GPUS-13
Nodes	4	8 + 1 head node	16 + 1 head node
GPU	8 M2090	16 M2090	32 M2090
CPU	8 X5670	16 X5670	32 X5670
Memory	24GB/node	48GB/node	48GB/node
Network	QDR InfiniBand	QDR InfiniBand	QDR InfiniBand
Rack	14U	14U	42U

http://www.supermicro.com/products/nfo/GPU.cfm?show=GPU_cluster

SUPERMICRO X9 (E5 SANDYBRIDGE) DUAL PROCESSOR GPU SERVER: SYS-2027GR-TRF

- Motherboard: **X9DRG-HF**
- 2U Chassis: **218G-R1800B**

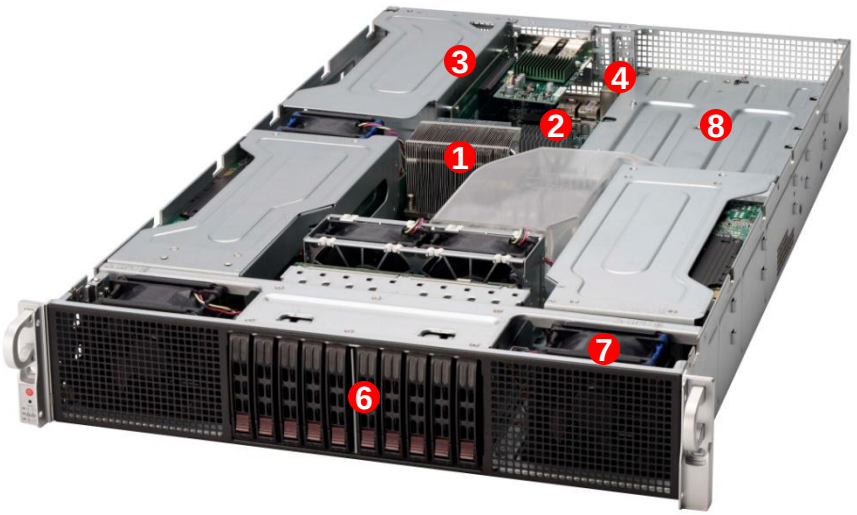

* Picture is reference only

Key features

- **Support 4 double width GPU cards. Optional riser card for 2 additional GPU cards.**
- **Platinum level redundant power supply**
- **10 x 2.5 inch HDD bays**
- **Smart server management tool**

1	**Processor Support** Dual Intel® Sandy Bridge EP (Socket R) series CPU
2	**Memory Capacity** 8 DIMM, Max of 256GB Reg. ECC DDR3 up to 1600MHz
3	**Expansion Slots** 4 PCI-e x16 Gen 3 for double width GPU cards -FM409: 4 NVIDIA M2090 GPU cards installed -FM475: 4 NVIDIA M2075 GPU cards installed 1 x4 STD size card 1 x8 LP card
4	**I/O ports** 1 VGA, 1 COM, 2 Gbit LAN, 2 USB 2.0, and 1 IPMI dedicated LAN port.
5	**System management** On board BMC (Baseboard Management Controllers) supports IPMI2.0, media/KVM over LAN. (Dedicated LAN port for management)
6	**Drive Bays** 10 hot-swap 2.5” drives bays
7	**System Cooling** 5 heavy duty 8cm fans with optimal fan speed control 1 air shroud
8	**Power Supply** 1800W Platinum level efficiency redundant power supply
	Dimensions H 3.5” (89mm) x W 17.2” (437mm) x D 30.5” (775mm)

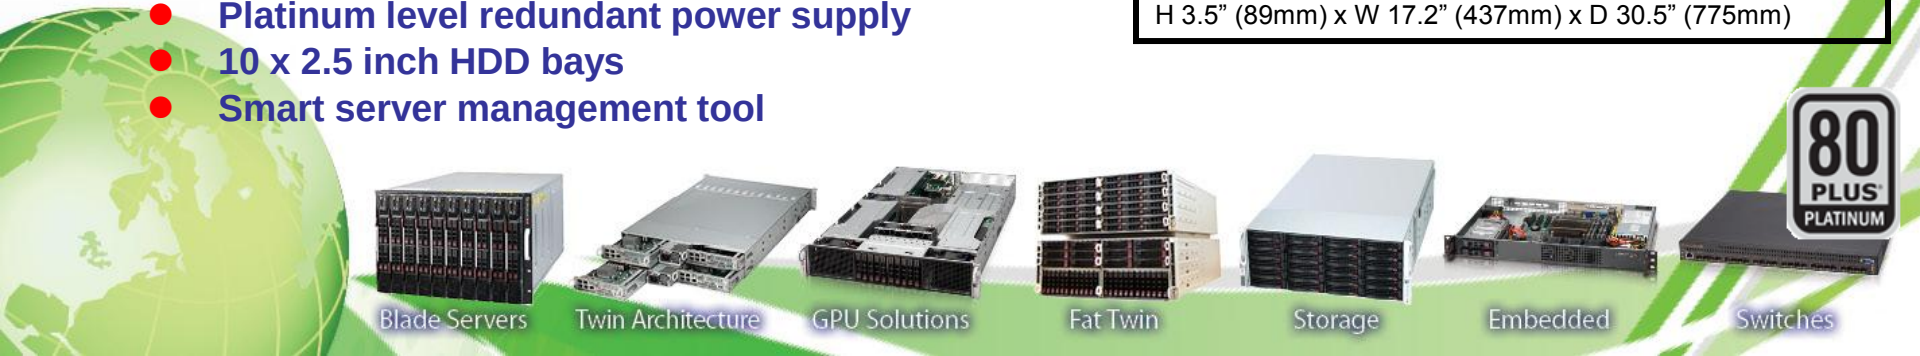

SUPERMICRO

DP GPU server: SYS-1027GR-TRF

- Motherboard: **X9DRG-HF**
- 1U Chassis: **118GQ-R1800B**

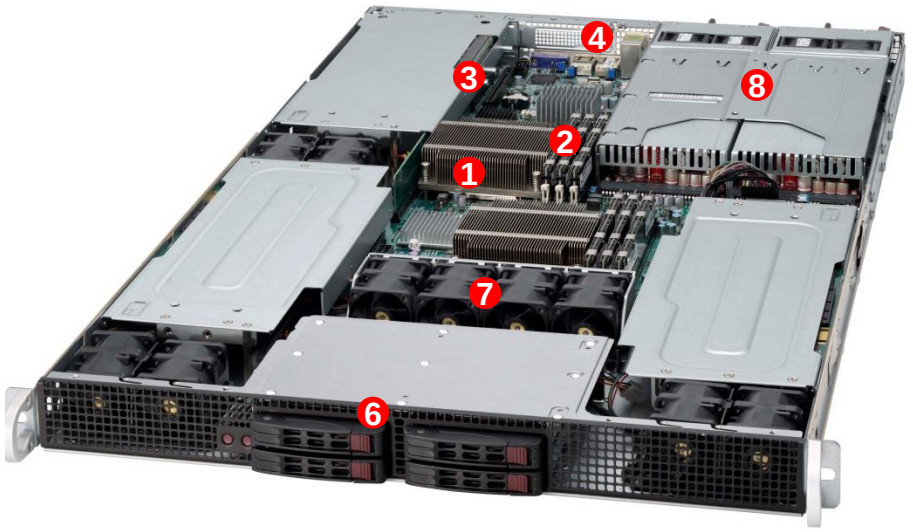

* Picture is reference only

1	**Processor Support** Dual Intel® Sandy Bridge EP (Socket R) series CPU
2	**Memory Capacity** 8 DIMM, Max of 256GB Reg. ECC DDR3 up to 1600MHz
3	**Expansion Slots** 3 PCI-e x16 Gen 3 for double width GPU cards -FM309: 3 NVIDIA M2090 GPU cards installed -FM375: 3 NVIDIA M2075 GPU cards installed 1 x8 (in x16 slot) LP card
4	**I/O ports** 1 VGA, 1 COM, 2 Gbit LAN, 2 USB 2.0, and 1 IPMI dedicated LAN port.
5	**System management** On board BMC (Baseboard Management Controllers) supports IPMI2.0, media/KVM over LAN. (Dedicated LAN port for management)
6	**Drive Bays** 4 hot-swap 2.5” drives bays
7	**System Cooling** 10 counter rotating fans with optimal fan speed control 1 air shroud
8	**Power Supply** 1800W Platinum level efficiency redundant power supply
	Dimensions H 1.7” (43mm) x W 17.2” (437mm) x D 30.6” (777mm)

Key features

- **Support up to 3 double width GPU cards**
- **Platinum level redundant power supply**
- **4 x 2.5 inch HDD bays**
- **Smart server management tool**

SUPERMICRO

UP GPU server: SYS-1017GR-TF

- Motherboard: **X9SRG-F**
- 1U GPU Chassis: CSE-118G-1200BP
- PWS: PWS-1K43F-1R

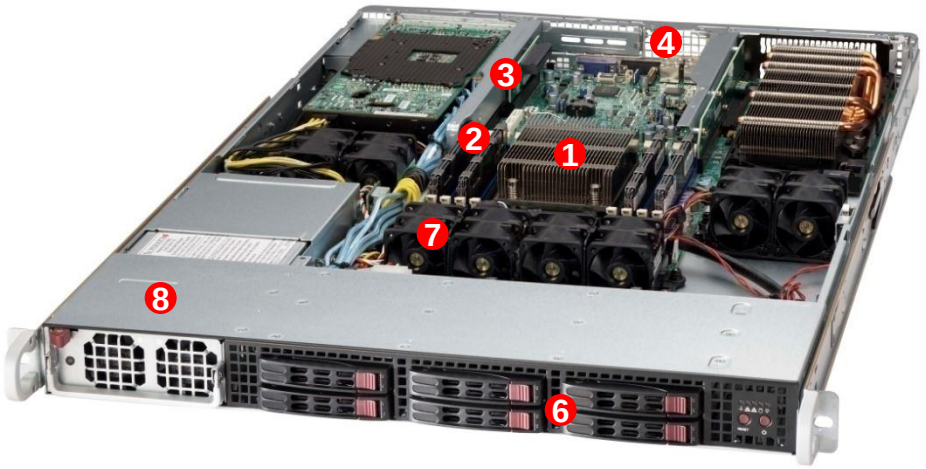

* Picture is reference only

Key features

- Sandy Bridge Romley Platform
- GPU optimized UP server
- Capable to support 2 x M series GPU (x16 slots)
- 1400W Platinum level super high efficiency power

1	**Processor Support** Single Intel® Sandy Bridge EP (Socket R) series CPU
2	**Memory Capacity** 8 DIMM, Max of 256GB Reg. ECC DDR3 or 64GB Un-buffered ECC DDR3 up to 1600MHz
3	**Expansion Slots** 2 PCI-e x16 Gen 3 for double width GPU cards -FM209: 2 NVIDIA M2090 GPU cards installed -FM275: 2 NVIDIA M2075 GPU cards installed 1 x8 Gen 3 LP card
4	**I/O ports** 1 VGA, 1 COM, 2 Gbit LAN (Intel dual Gbit LAN by Powervill), 2 USB 2.0, and 1 IPMI dedicated LAN port.
5	**System management** On board BMC (Baseboard Management Controllers) supports IPMI2.0, media/KVM over LAN. (Dedicated LAN port for management)
6	**Drive Bays** 6 hot-swap 2.5” drives bays
7	**System Cooling** 8 counter rotating fans w/ optimal fan speed control 3 air shroud
8	**Power Supply** 1400W Platinum level efficiency redundant power supply
	Dimensions H 1.7” (43mm) x W 17.2” (437mm) x D 28.2” (716mm)

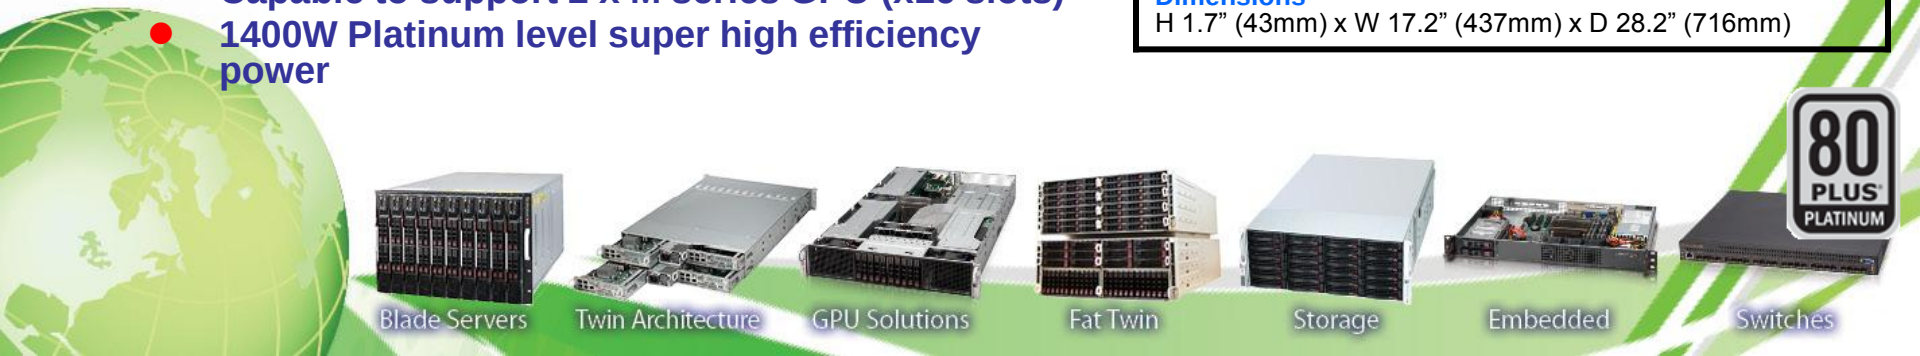

GPU Workstation: 7047GR-TRF / 7047GR-TPRF

Passive Cooling

- Motherboard: **X9DRG-QF**
- Tower/4U Chassis: CSE747TQ-R1620B

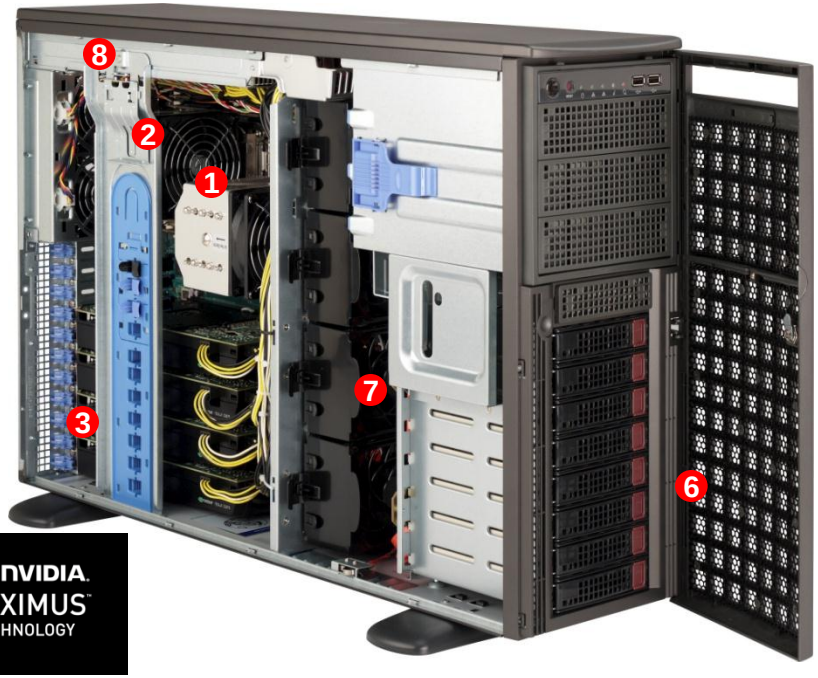

* Picture is reference only

1	**Processor Support** Dual Intel® Sandy Bridge EP (Socket R) series CPU
2	**Memory Capacity** 16 DIMM, 512GB DDR3 1600 MHz Reg. ECC
3	**Expansion Slots** 4 PCI-e 3.0 x16 for double width GPU cards 2 PCI-e 3.0 x8 (1 in x16 slot) 1 PCI-e 2.0 x4 (in x8 slot)
4	**I/O ports** 1 VGA, 2 COM, 2 Gbit LAN, 10 USB 2.0, and 1 IPMI dedicated LAN port.
5	**System management** On board BMC (Baseboard Management Controllers) supports IPMI2.0, media/KVM over LAN with dedicated LAN for system management.
6	**Drive Bays** 8 hot-swap 3.5” drives bay + 3 fixed 5.25” drives bay + 1 fixed 3.5” drive bay
7	**System Cooling** 4 heavy duty fans, 2 exhaust fans (4 for TPRF model), and 2 active heatsink w/ Optimal Fan Speed Control
8	**Power Supply** 1620W Platinum level efficiency redundant power supply
	Dimension H 17.8” (452mm) x W 7” (178mm) x D 26.5” (673mm)

NVIDIA MAXIMUS TECHNOLOGY

Key features

- **Supports 4 double width GPU cards + 3 PCIe cards**
- **Support up to 150W CPU**
- **1600W Platinum level redundant power supply**

Supermicro Power Design Makes Servers More Efficient

- Poor power efficiency increases power / cooling costs
  - Every watt wasted adds approx $1 in power costs and $1 in cooling costs.
- Traditional analog designs are now in the mid-80% range and reach diminishing returns in the low-90%'s.
  - Today's Supermicro analog designs are typically in the 92%-94%+ range
- Digital Power can achieve 95% and beyond.
- This translates to hundreds of dollars per node, thousands of dollars per rack and hundreds of thousands of dollars saved in a large data center.

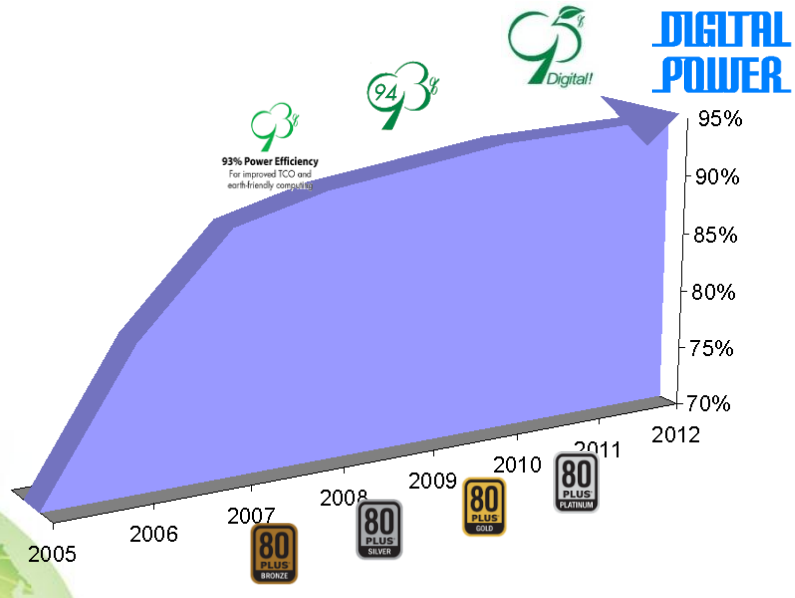

Efficiency

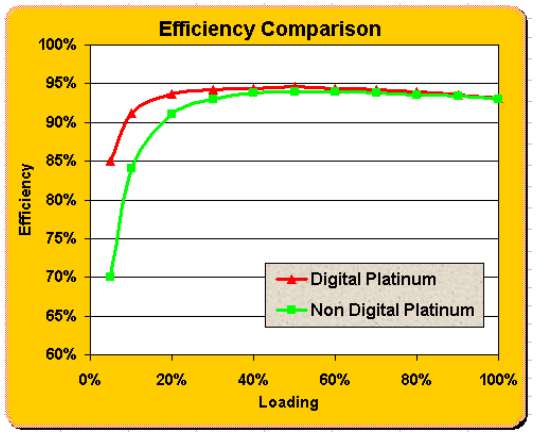

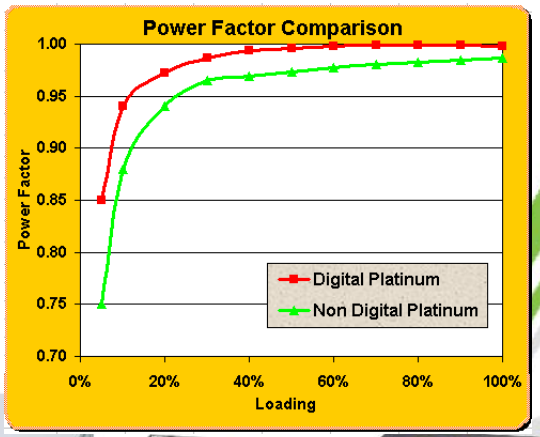

Feature Advantages: Quick Summary

- **Compute Density:**
  - Up to 4 Kepler / Tesla GPUs per Rack Unit; up to 4x advantage vs. general servers
- **Peak Performance:**
  - Scale-Up: Up to 5 GPUs per 4U workstation; 4 GPUs per 4-Way PSC platform
  - Fastest I/O: PCIe 3.0, FDR Infiniband, 10GbE, 6Gbs SAS
- **Scalability:**
  - Scale-Out in units of 1U, 2U, 4U, 7U; Supermicro Turn-key GPU Simclusters
  - Best Performance per watt per Dollar with Supermicro GPU Blades and GPU Twin architecture
- **Power Efficiency:**
  - Industry Leading Platinum Level. Analog 94%+ / Digital 95%+
- **Manageability:**
  - Supermicro IPMI system monitoring and Enterprise Server Management software suite (not covered in this presentation)
- **Reliability:**
  - Server Grade Design and Components, Optimized Thermals, Power Redundancy, Remote Management
- **Cost Effectiveness**
  - See your Supermicro Sales Representative

Recap

- **Supermicro Product Leadership**

➔ **Are the product benefits measurable?**

- Which specific products are GPU optimized?
- What are specific examples of product leadership?

Section

Q & A

- **Customer Endorsements**

➔ **Who are some of Supermicro's partners?**
- Why have they chosen Supermicro?

Global Deployment

Customer Endorsements

Global Customer Deployment

NASA

BAE SYSTEMS

NVIDIA

BR

PETROBRAS

peer1 hosting

SIEMENS

amazon.com

VT

ebay

TOTAL

Blade Servers | Twin Architecture | GPU Solutions | Fat Twin | Storage | Embedded | Switches

Typical Markets Served with GP / GPU Servers & Workstations

- Oil & Gas Exploration
- Medical Imaging
- Financial Modeling & Risk Analysis
- Weather Modeling
- Space Exploration
- Real-Time Transaction Security / Fraud Analysis
- Bio-Informatics & Life Sciences
- CAD / CAM Digital Content Creation Acceleration

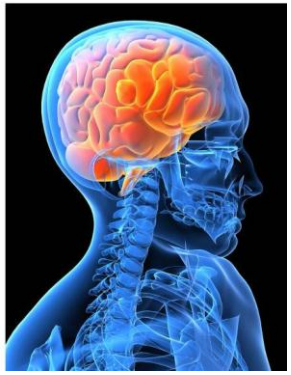


Blade Servers


Twin Architecture


GPU Solutions


Fat Twin


Storage


Embedded


Switches

Customer Responses

"…no other supplier was able to achieve the GPU density that I needed…"

"…this is the fastest system we've ever deployed."

"…Supermicro is consistently first-to-market."

"We have benefitted from direct access to product managers. We even have access to the original board and system designers!"

"With so MANY products to choose from, we really benefited from close interaction with our Supermicro sales rep."

"Supermicro… The fastest 1U-GPU system on the planet."

- Global Operations
- 24 / 7 / 365 Support

Summary

- Supermicro is growing fast to serve its customers
- Supermicro has the industry's broadest and most efficient Application Optimized system solutions
- Supermicro Keep IT Green™

BACKUP CONTENT

GTC Session

SUPERMICRO®

OPTIMAL SYSTEM ARCHITECTURE FOR GPU ACCELERATION

Model	 **SuperServer SYS-2027GR-TRF**	 **SuperServer SYS-1027GR-TRF**	 **SuperServer SYS-5017GR-TF**	 **SuperServer SYS-1017GR-TF**	 **SuperServer SYS-7047GR-TRF**
CPU	Sandy bridge DP	Sandy bridge DP	Sandy bridge UP	Sandy bridge UP	Sandy bridge DP
Memory	Up to 256GB in 8 DIMM	Up to 256GB in 8 DIMM	Up to 256GB in 8 DIMM	Up to 256GB in 8 DIMM	Up to 512GB in 16 DIMM
HDD	10 x 2.5”	4 x 2.5”	3 x 3.5”	6 x 2.5”	(8 + 2) x 3.5"
Expansion slot	4 GPU 1 PCI-e 3.0 x4 1 PCI-e 3.0 x8 (LP)	3 GPU 1 PCI-e 3.0 x8 (LP)	2 GPU 1 PCI-e 3.0 x8 (LP)	2 GPU 1 PCI-e 3.0 x8 (LP)	4 GPU 2 PCI-e 3.0 x8
Power Supply	1800W platinum redundant	1800W platinum redundant	1400W platinum	1400W platinum	1620W platinum redundant

* Pictures are reference only

SUPERMICRO

UP GPU server: SYS-5017GR-TF

- Motherboard: **X9SRG-F**
- 1U GPU Chassis: CSE-818G-1200BP
- PWS: PWS-1K43F-1R

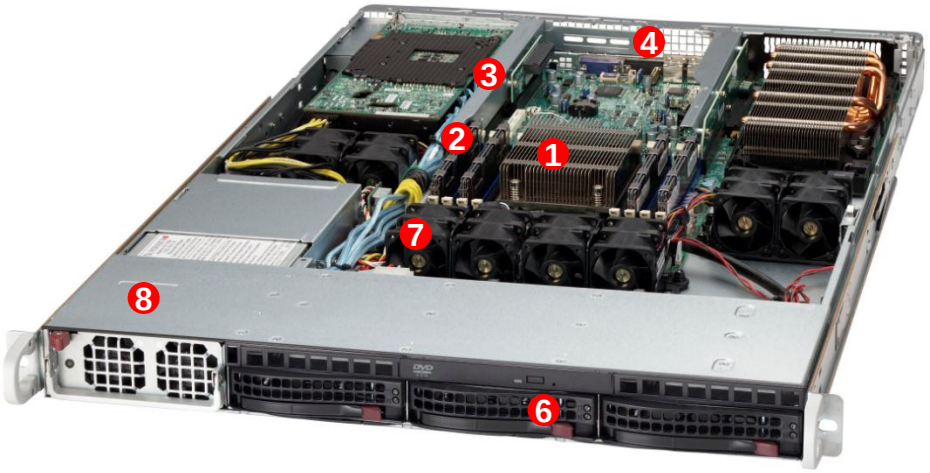

* Picture is reference only

Key features

- Sandy Bridge Romley Platform
- GPU optimized UP server
- Capable to support 2 x M series GPU (x16 slots)
- 1400W Platinum level super high efficiency power

1	**Processor Support** Single Intel® Sandy Bridge EP (Socket R) series CPU
2	**Memory Capacity** 8 DIMM, Max of 256GB Reg. ECC DDR3 or 64GB Un-buffered ECC DDR3 up to 1600MHz
3	**Expansion Slots** 2 PCI-e x16 Gen 3 for double width GPU cards -FM209: 2 NVIDIA M2090 GPU cards installed -FM275: 2 NVIDIA M2075 GPU cards installed 1 x8 Gen 3 LP card
4	**I/O ports** 1 VGA, 1 COM, 2 Gbit LAN (Intel dual Gbit LAN by Powervill), 2 USB 2.0, and 1 IPMI dedicated LAN port.
5	**System management** On board BMC (Baseboard Management Controllers) supports IPMI2.0, media/KVM over LAN. (Dedicated LAN port for management)
6	**Drive Bays** 3 hot-swap 3.5” drives bays
7	**System Cooling** 8 counter rotating fans w/ optimal fan speed control 3 air shroud
8	**Power Supply** 1400W Platinum level efficiency redundant power supply
	Dimensions H 1.7” (43mm) x W 17.2” (437mm) x D 28.2” (716mm)

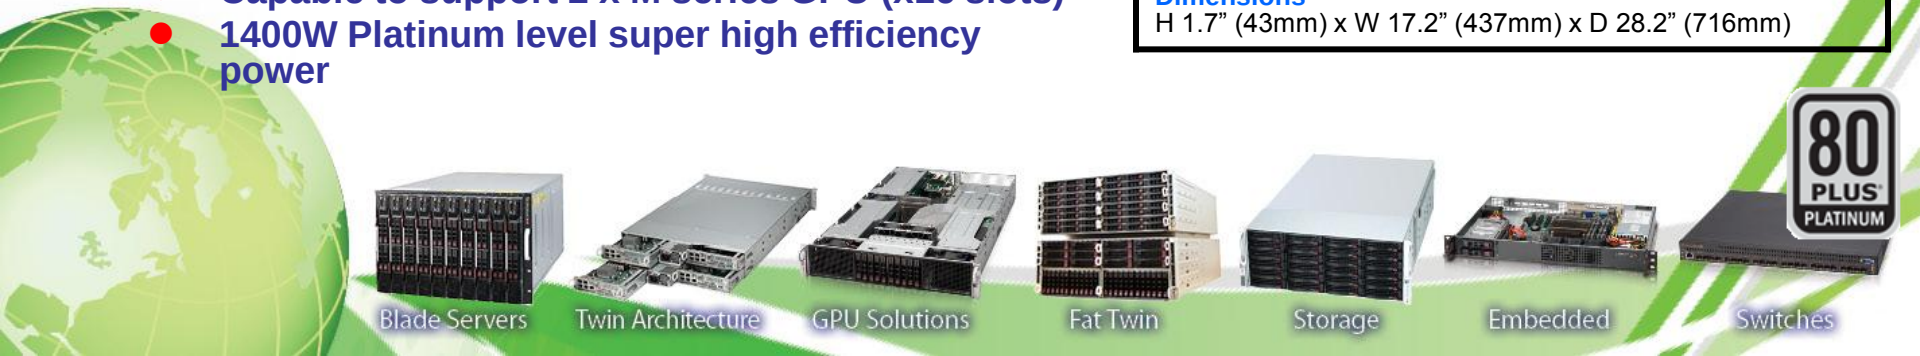

Love Our Mother Earth

- **High efficiency systems in computing could help alleviate the world's biggest environmental challenges while driving profitable growth for users**
- **Compared to typical server systems:**
  - **Each 94%+ high efficiency power server can save up to ~$100 USD per node per year in energy**
  - **986 tons of CO_2 can be reduced from the atmosphere for every 1000 high efficiency system deployment**
  - **Worldwide adoption of high-efficiency solutions could save the world ~$9 billion USD per year in electricity or ~1.2 billion trees for the sequestration of CO_2**

Green =

We Keep IT Green™