Specification

Tesla S870
GPU Computing System

March 27, 2008 | SP-03685-001_v01c

Document Change History

Version	Date	Responsible	Description of Change
01b			Initial release
01c	March 27, 2008	Geoff Ballew	Updated global part number and improved system architecture diagrams.

Table of Contents

Tesla S870 Overview 1
Key Specifications 1
Computing Processors 1
Mechanical Overview 1
Operating Environment 1
System Architecture 2
Configuration 4
Mechanical Specification 5
System Chassis 5
Host Adapter Card 6
PCI Express Cable 7
Thermal Specifications 8
Thermal Testing Summary 8
Support Information 10
Languages 10
Certificates and Agencies 10
Certifications 10
Agencies 10

List of Figures

Figure 1. Tesla S870 System Architecture 2
Figure 2. Tesla S870 Connected to a Single Host System 3
Figure 3. Tesla S870 Connected to Two Host Systems 3
Figure 4. System Chassis Drawing 5
Figure 5. Host Adapter Card with Bracket Options 6
Figure 6. PCI Express Cable 7

List of Tables

Table 1. System Configuration 4
Table 2. Test Parameter Highlights 8
Table 3. Thermal Results and Specification 8
Table 4. Environment Specifications and Conditions 9

Tesla S870 Overview

The NVIDIA® Tesla™ S870 GPU Computing System is a 1U rack-mount system with four Tesla graphics processing units (GPUs). This system connects to a host system via two PCI Express cables and host adapter cards that plug into a PCIe slot in the host system(s). The host adapter cards are compatible with both PCIe 1x and PCIe 2.0 systems.

Key Specifications

Computing Processors

- Four Tesla C870 GPUs
- 6.0 GB of high speed memory, configured as 1.5 GB for each GPU

Mechanical Overview

- Physical dimensions
  - System: 1.73 inches high × 17.5 inches wide × 31 inches deep
  - PCIe cable:
    - Standard: 0.5 meters in length
    - Optional: 2.0 meters in length
  - Two host adapter cards: PCI Express low profile form factor
- External Connectors
  - Two cable connectors for x16 PCI Express
  - C19 format female connector for power cord

Operating Environment

- Temperature: 5 °C to 35 °C
- Relative humidity: 10 % to 80 % non-condensing
- 115 cubic feet per minute airflow front to back

System Architecture

The Tesla S870 GPU computing system is based on the 8 series GPU from NVIDIA. It can be connected to a single host system via 2 PCI Express connections to that host, or connected to two separate host systems via 1 PCI Express connection to each host. Each PCIe connection leads to 2 of the 4 Tesla Processors. If only one PCIe connection is made, then only 2 of the 4 Tesla Processors will be active, but the system will work. The Tesla S870 system architecture is shown below in Figure 1.

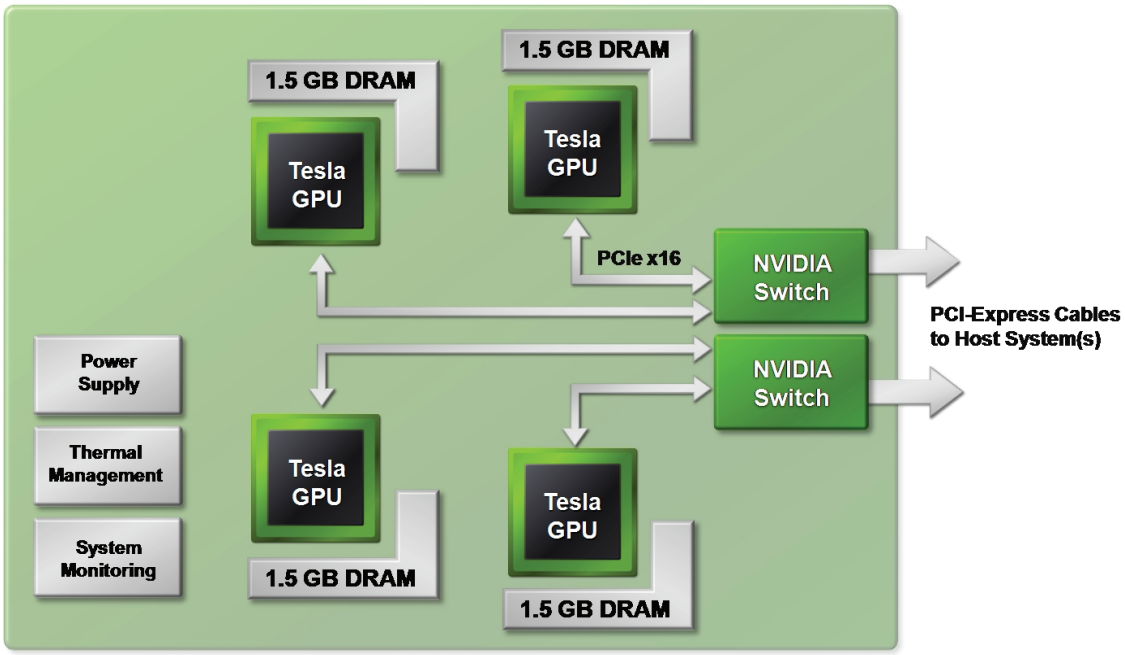

Figure 1. Tesla S870 System Architecture

Figure 2 shows how the Tesla S870 can be connected to a single host system with both PCIe connections routed to the host.

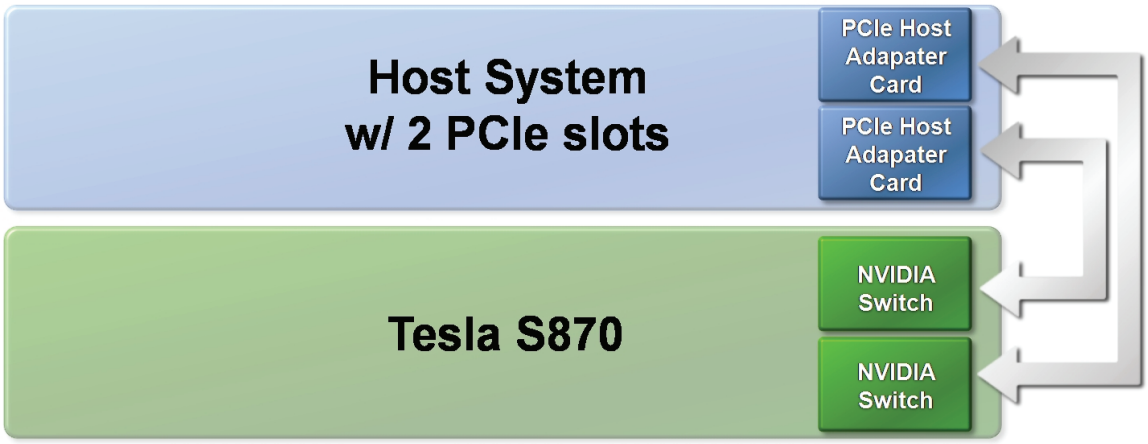

Figure 2. Tesla S870 Connected to a Single Host System

Figure 3 shows how the Tesla S870 can be connected to two host systems with one PCIe connection routed to each host system. This has the benefit of requiring only a single available PCIe slot in the host system.

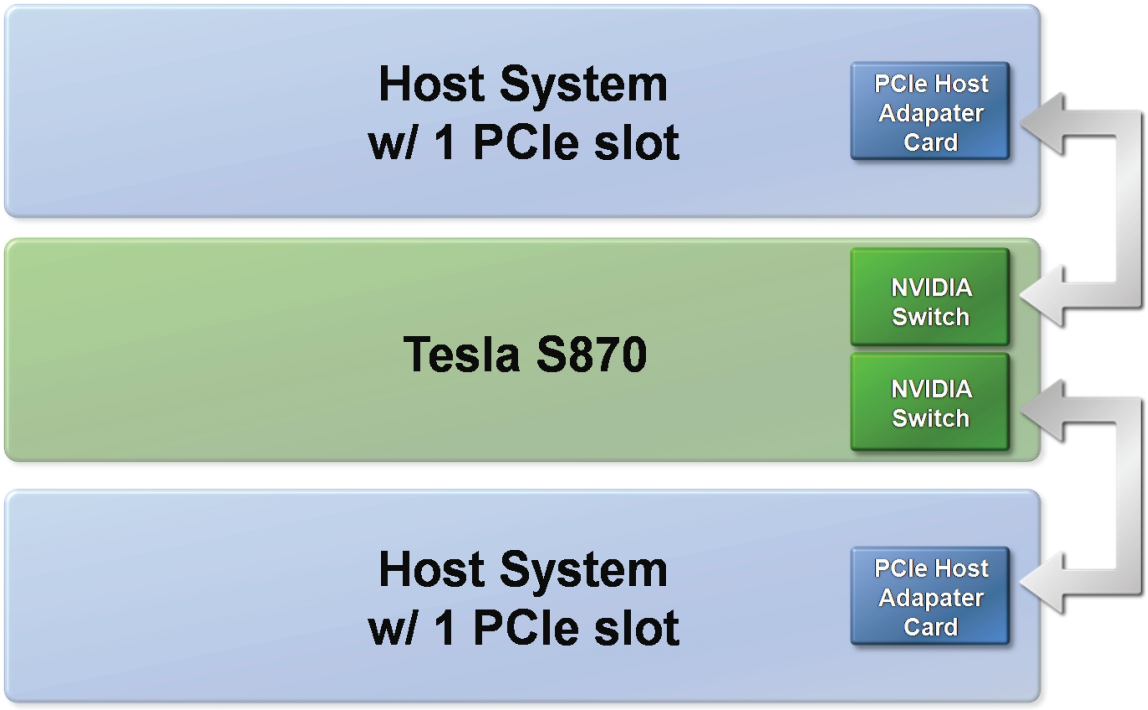

Figure 3. Tesla S870 GPU Connected to Two Host Systems

Configuration

There is one configuration available (Table 1) for the Tesla S870 computing system. Geographical SKUs have localized power cords but the system configuration is the same.

Table 1. System Configuration

Specification	Description
Ordering Part Number	920-20709-0001-000 (Global)
GPU	Tesla C870 GPU
GPU Processor clock	1350 MHz
GPU Memory clock	800 MHz
Memory configuration	6.0 GB total configured as 1.5 GB per GPU
Memory I/O	384-bit per GPU
System I/O	Two 2 x16 PCI Express cable connections. Each connects to 2 of the GPUs in the Tesla S870
PCI Express Cables	2 cables, 0.5 meter length each

System Chassis

The Tesla 2870 (Figure 2) is a 1U form factor chassis and conforms to the EIA 310E specification for 19-inch 4-post racks with 900 mm to 1000 mm depth. The chassis dimensions are 1.73 inches high × 17.5 inches wide × 31 inches deep.

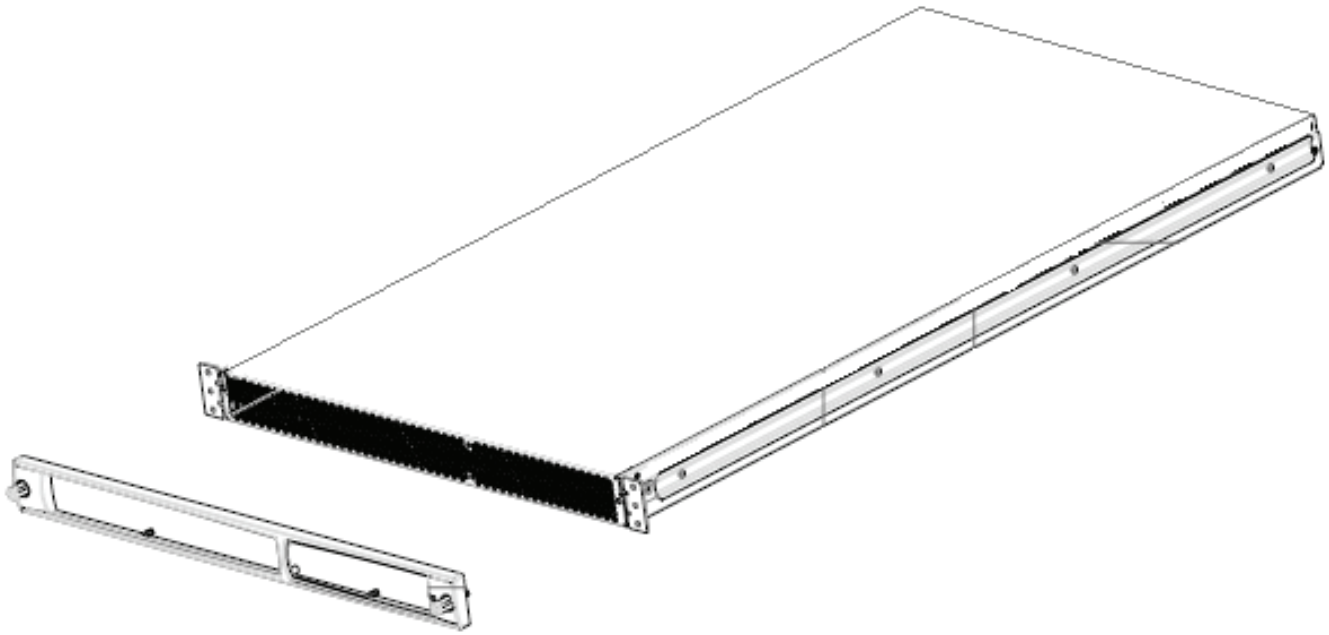

Figure 4. System Chassis Drawing

Host Adapter Card

The NVIDIA interface card (Figure 3) conforms to the PCI Express low profile form factor. This card is compatible with both PCI Express 1x and PCI Express 2.0 systems.

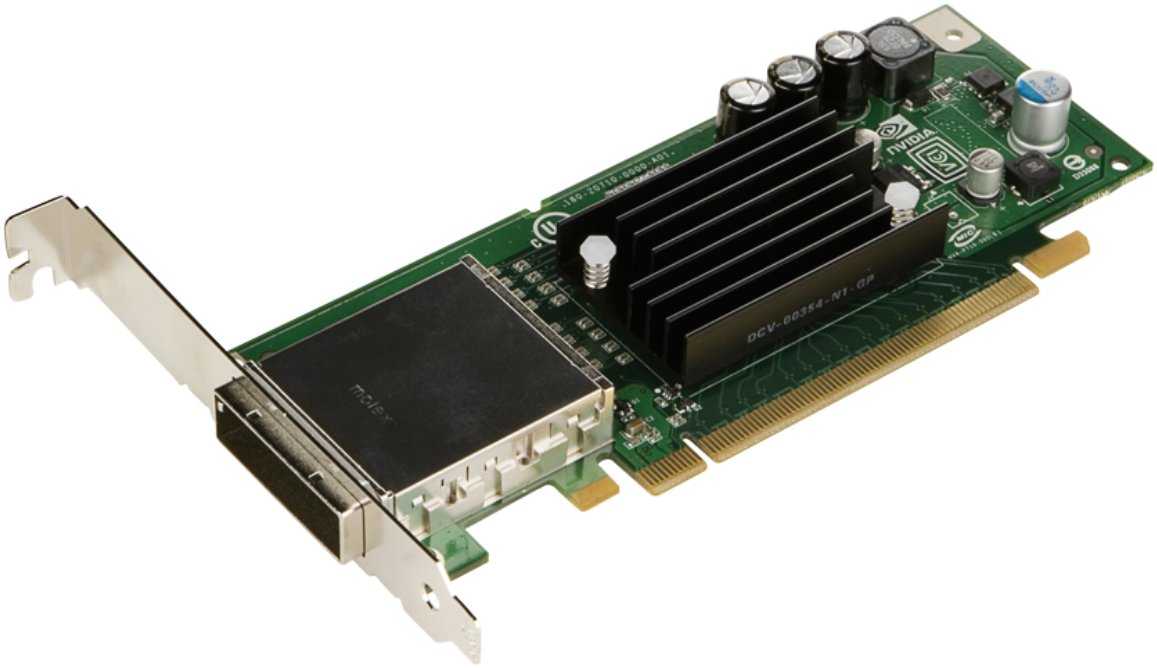

Figure 5. Host Adapter Card with Bracket Options

PCI Express Cable

The Tesla S870 includes two 0.5-meter PCI Express cables (Figure 4). These connect the Tesla S870 to the host system(s) using the host adapter cards. A 2.0-meter version of the cable is also available.

Note: For Figure 4 the dimensions are in millimeters unless otherwise labeled.

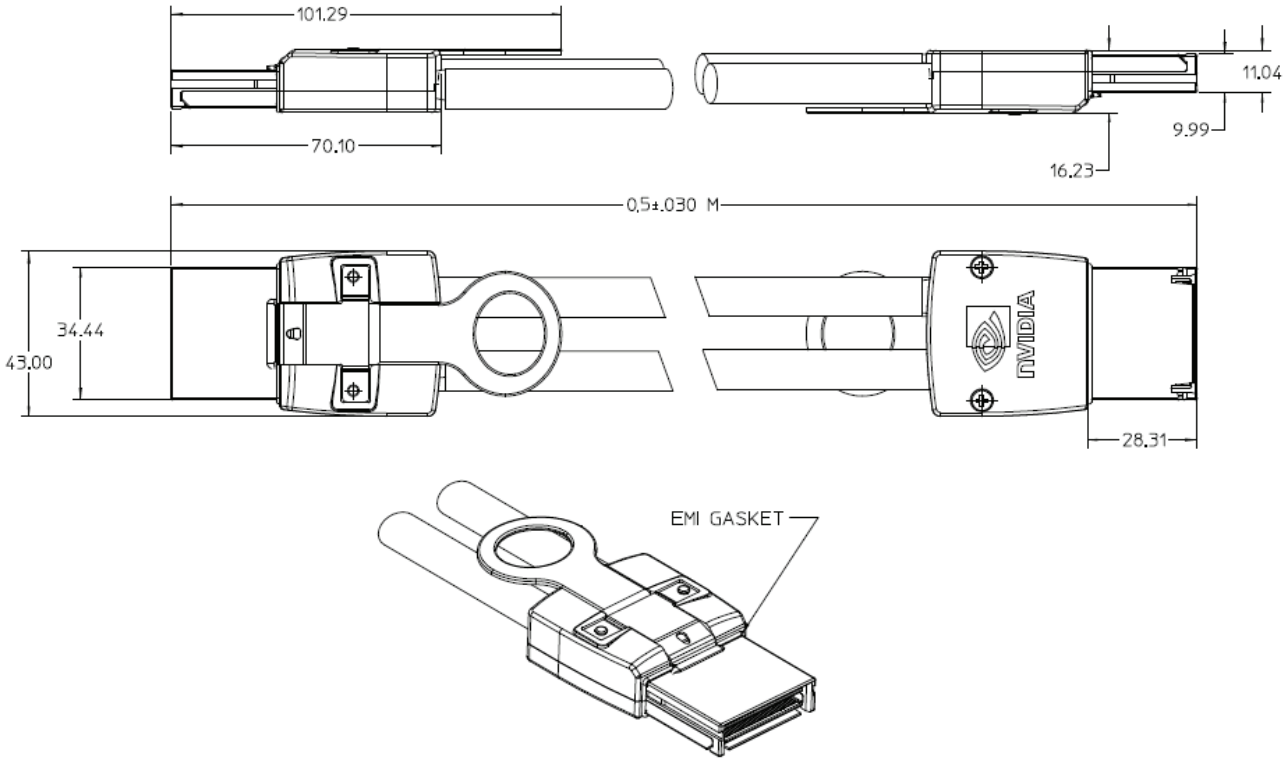

Figure 6. PCI Express Cable

Thermal Testing Summary

The thermal testing results are preliminary. Thermal testing will be detailed in a thermal qualification summary report.

Table 2. Test Parameter Highlights

Parameter	Value
Fan speed	50% duty cycle
Ambient temperature	35 °C
Duration	60 minutes
Components instrumented for the test	All 4 GPUs, 2 worst-case FETs Both PCIe switches
Test applications	NVIDIA thermal stress test NVIDIA binomial options test

Table 3. Thermal Results and Specification

Test Application	Average Power (Watts)	GPU Tjunction(°C)*	PCIe Switch Tjunction(°C)*	FET Tjunction(°C)*
NVIDIA thermal stress test	790	76	76	70

* Junction temperature as reported by NVIDIA thermal sensor. Only the highest value is reported for multiple chips. (e.g. There are four GPUs in the system, the reported value is the highest of the four.)

Table 4. Environment Specifications and Conditions

Specifications	Conditions
Operating temperature	5 °C to 35 °C
Operating humidity	10 to 80 % RH, non-condensing
Operating altitude	5000 feet mean sea level (MSL)
Operating shock	Half sine 40g, 2ms duration, 3 axis
Operating vibration	Sinusoidal 0.25g, 10 to 500 Hz, 3 axis. Random 1.0 Grms, 10 to 500 Hz
Acoustics	66 dBa at 1 meter in front of system
Non-operating temperature	0 °C to 65 °C

Languages

Language support for the Tesla S870 is English (U.S.) only at this time.

Certificates and Agencies

Certifications

There are no certifications planned at this time.

Agencies

- Australia Communication Authorities (C-Tick)
- Bureau of Standards, Metrology and Inspection (BSMI)
- Conformité Européenne (CE)
- Federal Communications Commission (FCC)
- Interference-Causing Equipment Standard (ICES)
- Ministry of Information and Communication (MIC)
- Underwriters Laboratories (UL)
- Voluntary Control Council for Interference (VCCI)

Notice

ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, “MATERIALS”) ARE BEING PROVIDED “AS IS.” NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE.

Information furnished is believed to be accurate and reliable. However, NVIDIA Corporation assumes no responsibility for the consequences of use of such information or for any infringement of patents or other rights of third parties that may result from its use. No license is granted by implication or otherwise under any patent or patent rights of NVIDIA Corporation. Specifications mentioned in this publication are subject to change without notice. This publication supersedes and replaces all information previously supplied. NVIDIA Corporation products are not authorized for use as critical components in life support devices or systems without express written approval of NVIDIA Corporation.

Trademarks

NVIDIA, the NVIDIA logo, and Tesla are trademarks or registered trademarks of NVIDIA Corporation. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2008 NVIDIA Corporation. All rights reserved.

NVIDIA Corporation | 2701 San Tomas Expressway | Santa Clara, CA 95050 | **www.nvidia.com**