



The Future
of
INDUSTRIAL COMPUTING

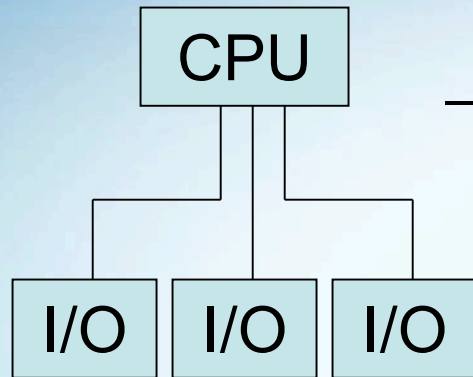
SuperSwitch & ExpressNet Technology Overview

July 2007

ONE STOP
SYSTEMS


PCI Express Basics

Two Architectures



- Tree – One CPU and multiple I/O boards
 - Available now

- Network – Multi CPUs, Multi I/O
 - Requires SuperSwitch H/W and ExpressNet_™ S/W
 - Available late-2007




PCI Express Basics Terminology

- Host = Upstream = Root Complex
 - Refers to CPU board in single-root system
- Target = Downstream = PCIe Endpoint
 - Refers to I/O board or intelligent I/O board

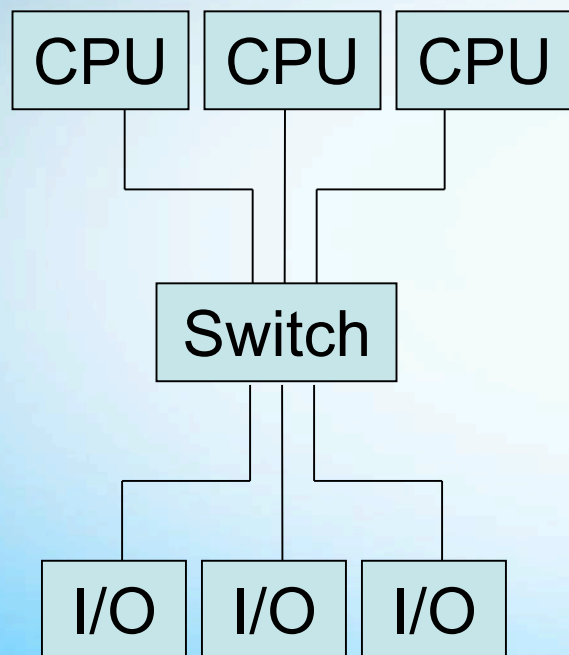
PCI Express Basics

Lanes – The Key to Performance




- Each lane consists of two 2.5Gbit/sec differential pairs
 - Separate Transmit and Receive pairs
 - Approximately 500MB/sec bandwidth total
- Multiple lanes are used to increase performance
 - x1 → 500MB/sec x1 (pronounced “by one”)
 - x4 → 2GB/sec
 - x8 → 4GB/sec
 - x16 → 8GB/sec
- Gen 2 clock rate will be 5Gbit/sec
 - Doubles bandwidth
 - Components will auto-detect if they can run at 5Gbits/sec
 - Products available in 2007 or 2008 timeframe

SuperSwitch Overview



- A way of using PCIe to communicate between multiple CPUs
 - PCIe as a networking architecture
 - High performance
 - Low cost
 - Based on industry standards
- Works for both cabled systems and backplane systems
- Both hardware and software components

SuperSwitch Technology





- One CPU initializes the switch components and handles the fabric management
 - The root complex CPU
- Non-transparent (NT) PCIe ports are used to make the other CPU elements PCIe end points
 - Each PCIe CPU looks like an intelligent I/O board to the system
- ExpressNet driver sets up memory windows from each CPU to all the other CPUs in the network
 - Driver allows direct DMA transfers
 - Driver implements TCP/IP interface
 - Looks like a NIC card to applications

Hardware Implementations

- Cabled systems with SuperSwitch™
 - Two system solution
 - 3 to 5-port solution
 - > 5-port solution
- Backplane systems with SuperSwitch™


PCIe-over-cable...Interface Products



ONE STOP
SYSTEMS

SuperSwitch 1

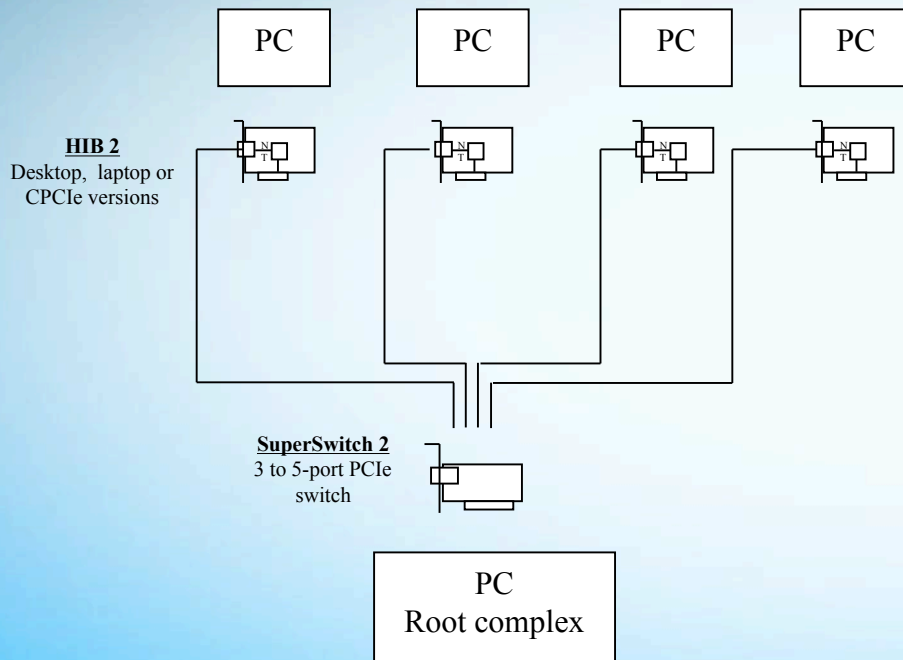
PCI Express-over-cable solution for point-to-point PC communications



- Uses HIB-2 in one PC and HIB-3 in the other
- Demonstrated performance of 725MB/sec
- Supported in Phase I of ExpressNet

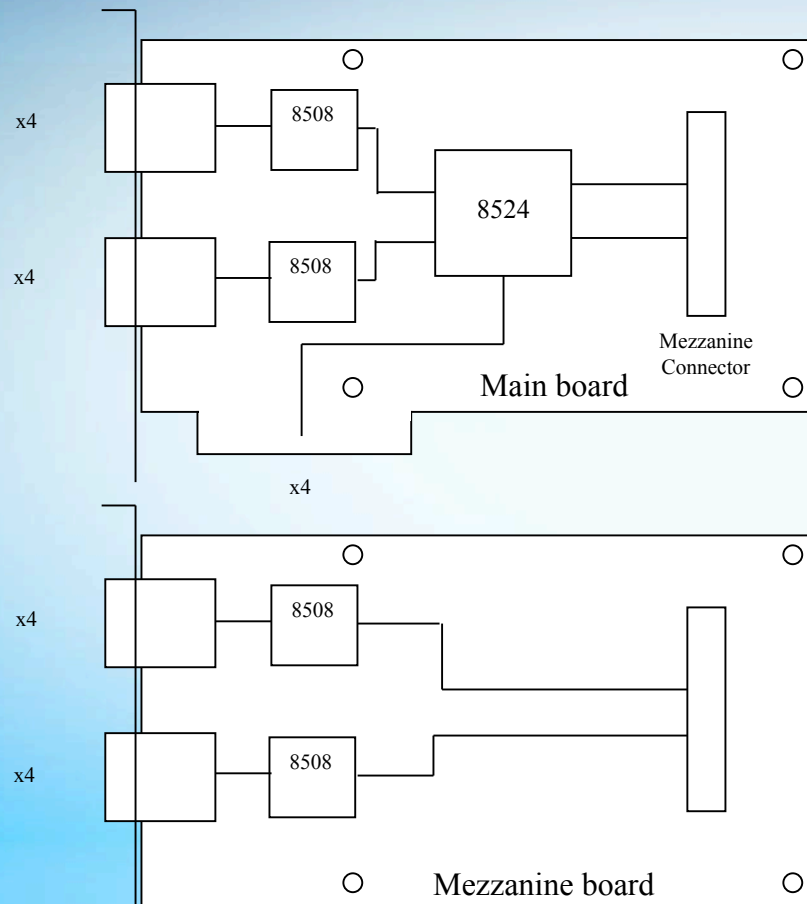
SuperSwitch 2

PCI Express-over-cable networking for 3 to 5 PCs



- Low-cost SuperSwitch
 - 3 or 5-port PCIe add-in board
 - 5-port version is two slots wide
 - Connects to 4 external PCs through PCIe x4 cable
- Standard HIB-2 boards used in node PCs
 - Desktop
 - Laptop
 - CPCle

SuperSwitch 2

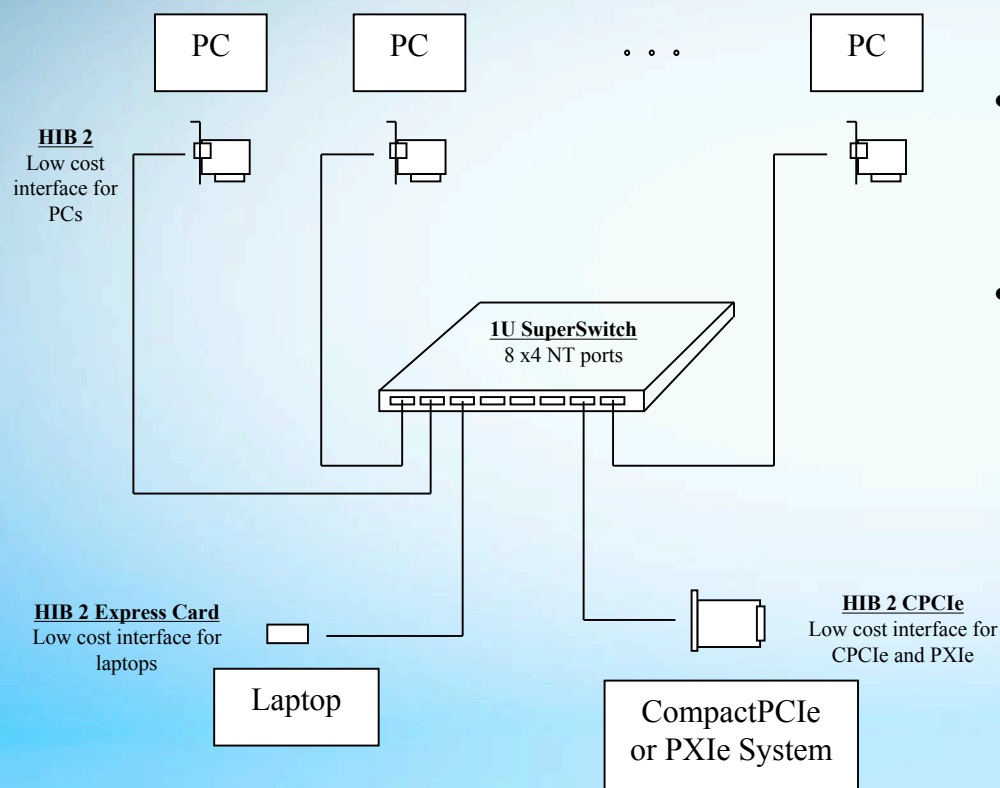


Low-cost 3 to 5-port:

- Two slot wide PCIe x4 board set
- Four external x4 connectors
- x8 Gen 2 version will follow

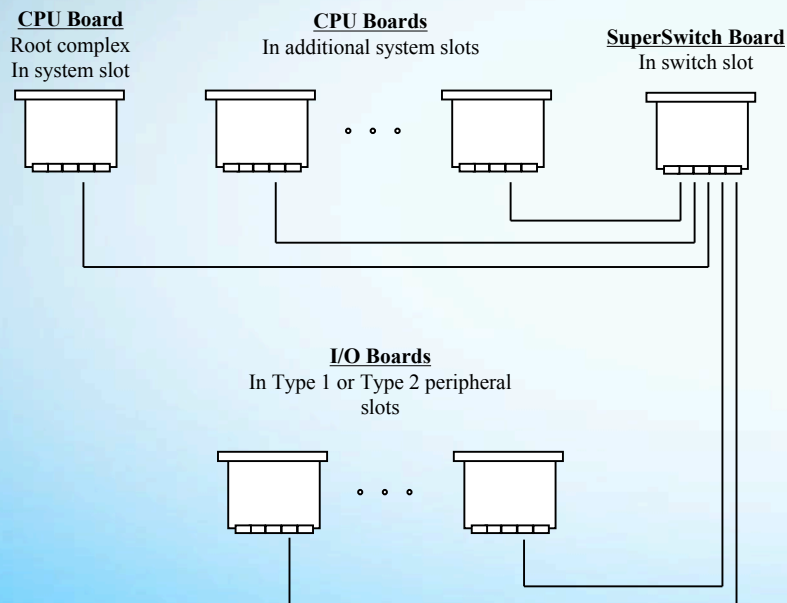
SuperSwitch 3

PCI Express-over-cable networking for 6 to 8 PCs



- Based on new 1U SuperSwitch™
 - NT ports in switch
 - Root complex CPU in switch
- Allows other CPUs to utilize generic PCIe over cable interface boards
 - Lower cost
 - Higher performance
 - Available in more form-factors

CompactPCIe Backplane System with SuperSwitch™

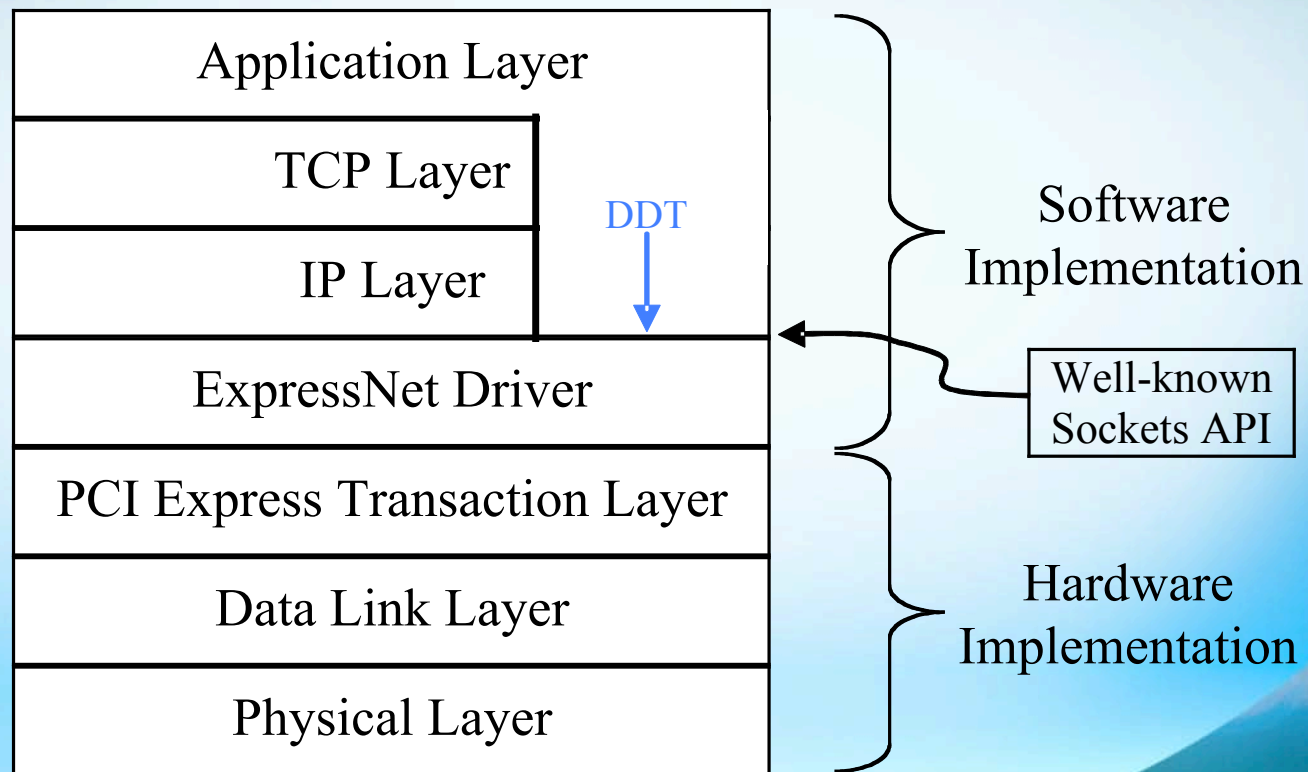


- Based on new CPCle SuperSwitch™
 - NT ports in switch
- Root complex CPU in separate system slot
- CPCle CPU and I/O boards are used unchanged
- Consolidates backplane interfaces into one unified architecture
 - PCIe used for both CPU-to-I/O and CPU-to-CPU communications
- SuperSwitch™ easily ported to other backplane form factors

ExpressNet Software

- Initialization and fabric management S/W
 - Runs on root complex CPU
 - Sets up network and manages hot-swap events
- ExpressNet Driver S/W
 - Runs on each peripheral CPU
 - Configures NT bridges
 - Two software interfaces
 - Direct data transfer
 - TCP/IP
- Demo program that verifies H/W setup and validates system performance


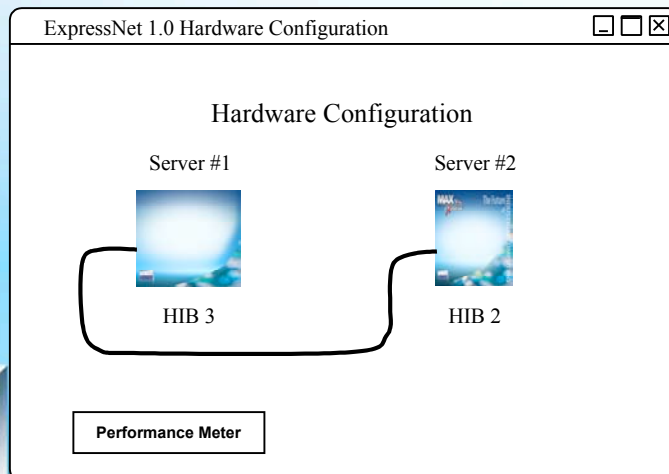
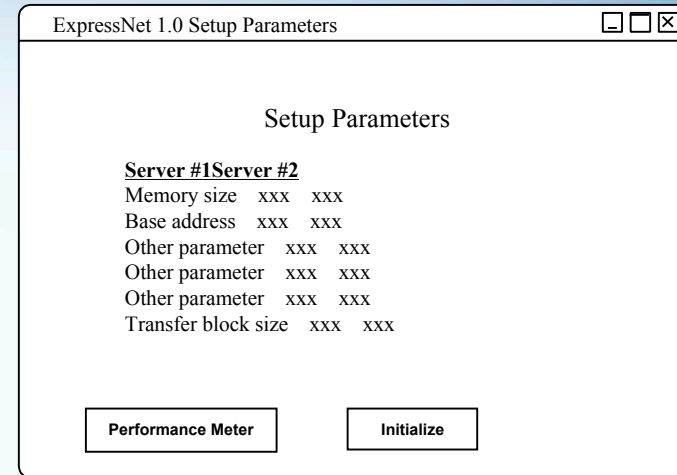
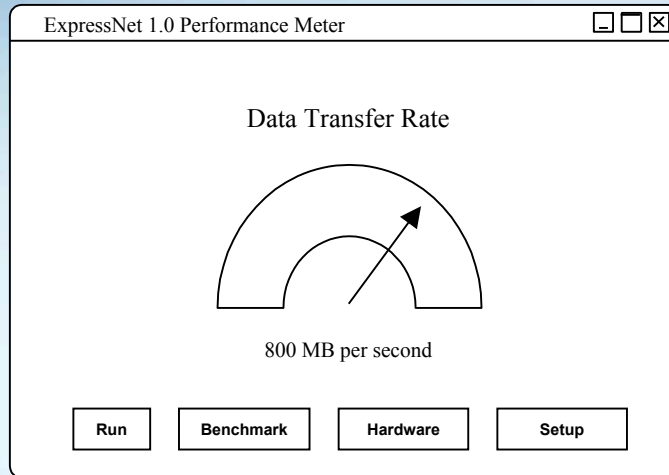
ExpressNet OSI Model



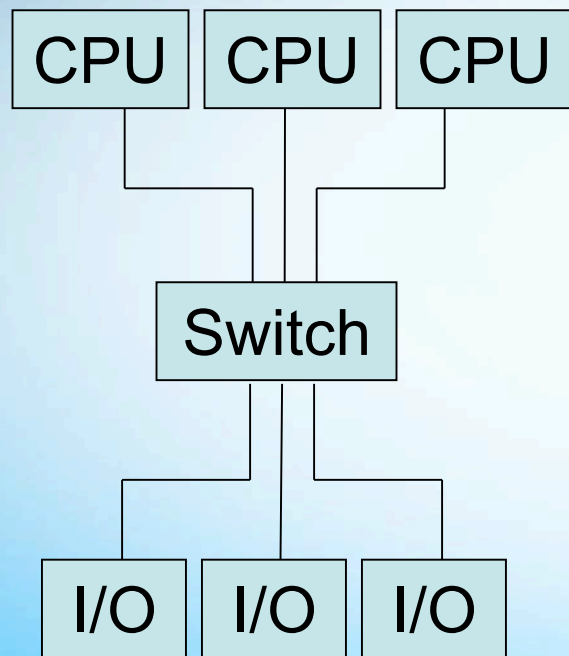
ExpressNet Release Phases

- Phase I
 - Supports 2 system solution
 - Linux OS only
 - Demonstration program
 - Direct data transfer support
- Phase II
 - Supports dual and multi-node solutions
 - DDT and TCP/IP support
 - Linux OS only
- Phase III
 - Supports dual and multi-node solutions
 - Linux and Windows Oss
 - Adds advanced features

Phase I ExpressNet Release



Summary

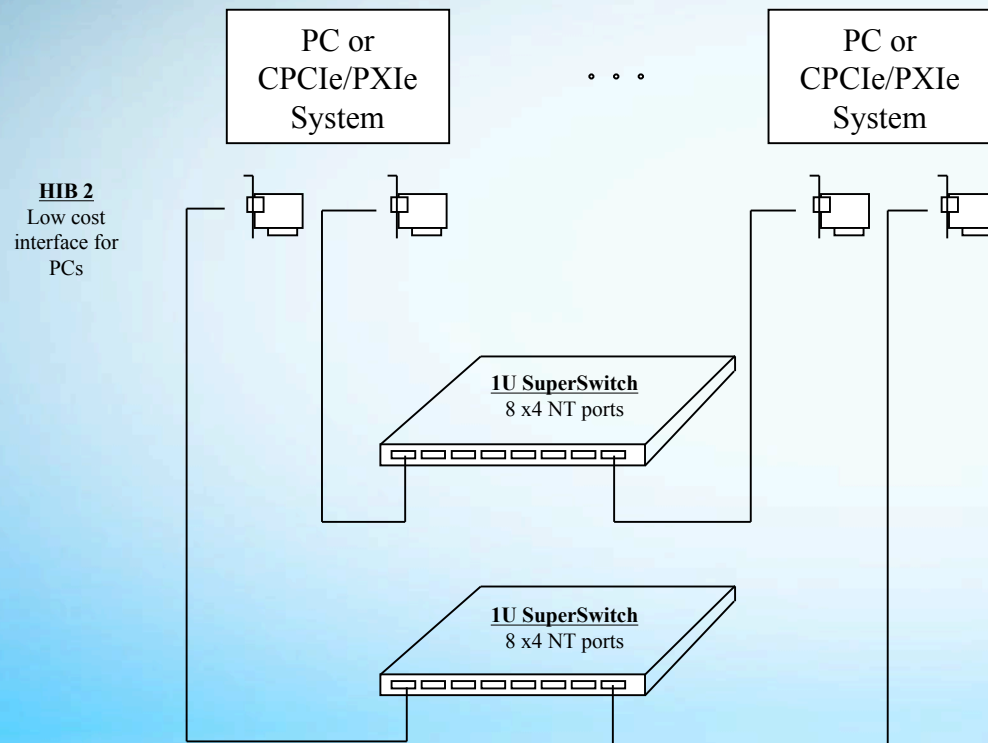


- SuperSwitch with ExpressNet driver allows multiple CPUs to communicate over PCIe
- High performance and low cost
- Runs over cabled and backplane systems
 - Provides bus unification in backplane systems
- Supports hot-swap and system redundancy

Hot Swap Capabilities

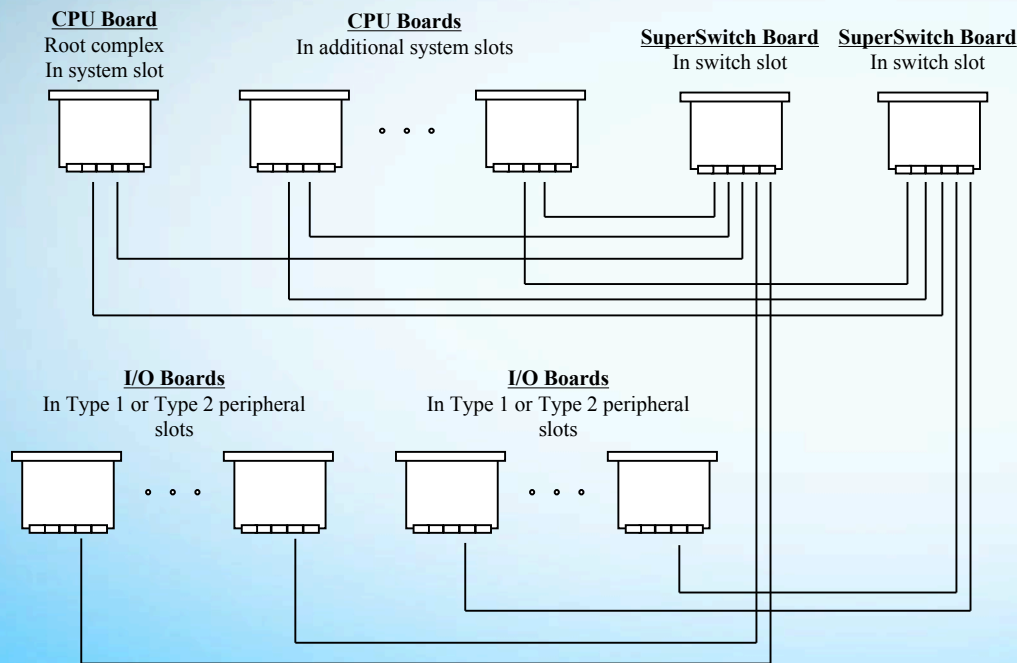
- In cabled systems
 - Any node can be connected or disconnected while the other nodes are running
- In backplane systems
 - Any board may be inserted or removed while the system is running
- Fabric manager software running on root complex CPU
 - Detects insertion or removal of new hardware
 - Notifies nodes of fabric changes
 - Reconfigures NT ports as needed

Redundant Cabled System Architecture



- Two SuperSwitch 1U units can be used to provide redundant switching
- Each PC or CompactPCIe system needs to include two PCIe over cable interface boards

Redundant Backplane System Architecture



- Two SuperSwitch CPCle switch boards can be used to provide redundant switching
- All CPCle CPU boards include multiple PCIe ports
- Most off-the-shelf I/O boards support only a single PCIe port so can only be connected to a single switch
 - Custom I/O boards could be designed with two PCIe ports
- The CPCle backplane needs to be designed to support this architecture

ExpressNet Software Enables Direct Data Transfer

- Low latency, low overhead protocol
- Well-known sockets API
- Simple protocol to move data

ExpressNet Software – TCP/IP

- TCP/IP fits above the ExpressNet driver
- Convenient, interface looks like NIC card
- Enables many network features

ExpressNet Software – Details

- Uses linked chain of buffers
- Descriptors define buffers and data structures
- Polled or interrupt driven
- Primarily uses PCI writes for efficiency