HyperTransport™ *HTX™*

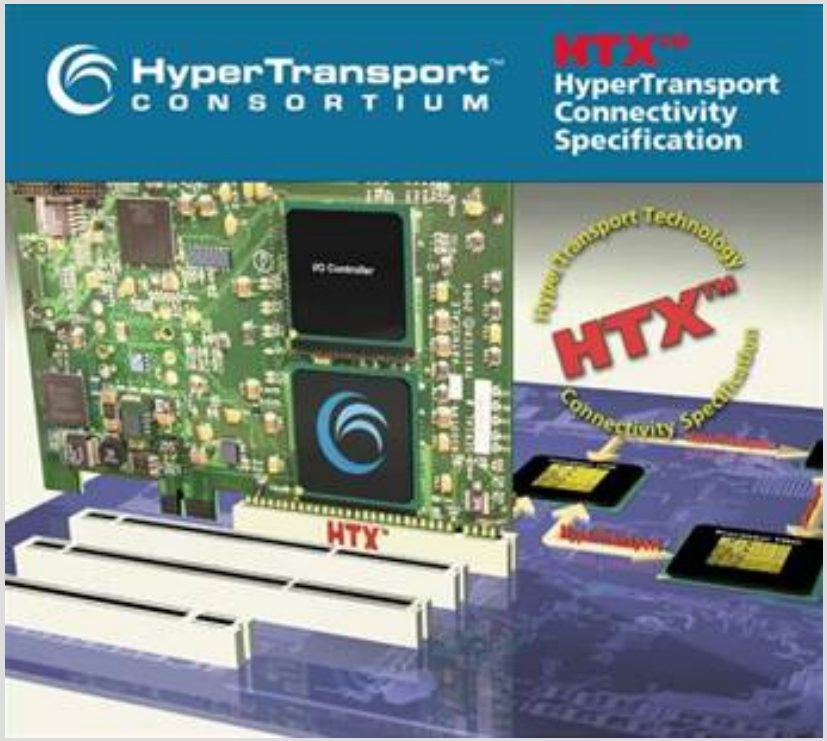

- Market Pull
- Technology Evolution
- Core Values
- Main Applications
- Products Ecosystem
- Development Support
- BIOS Support
- Why HTX

Market Pull

HyperTransport™ CONSORTIUM

High-Performance Computing Has Overflowed Original Boundaries

"HPC is now pervasive (users not specialists)"
Airbus Industries - ISC 2006, Dresden

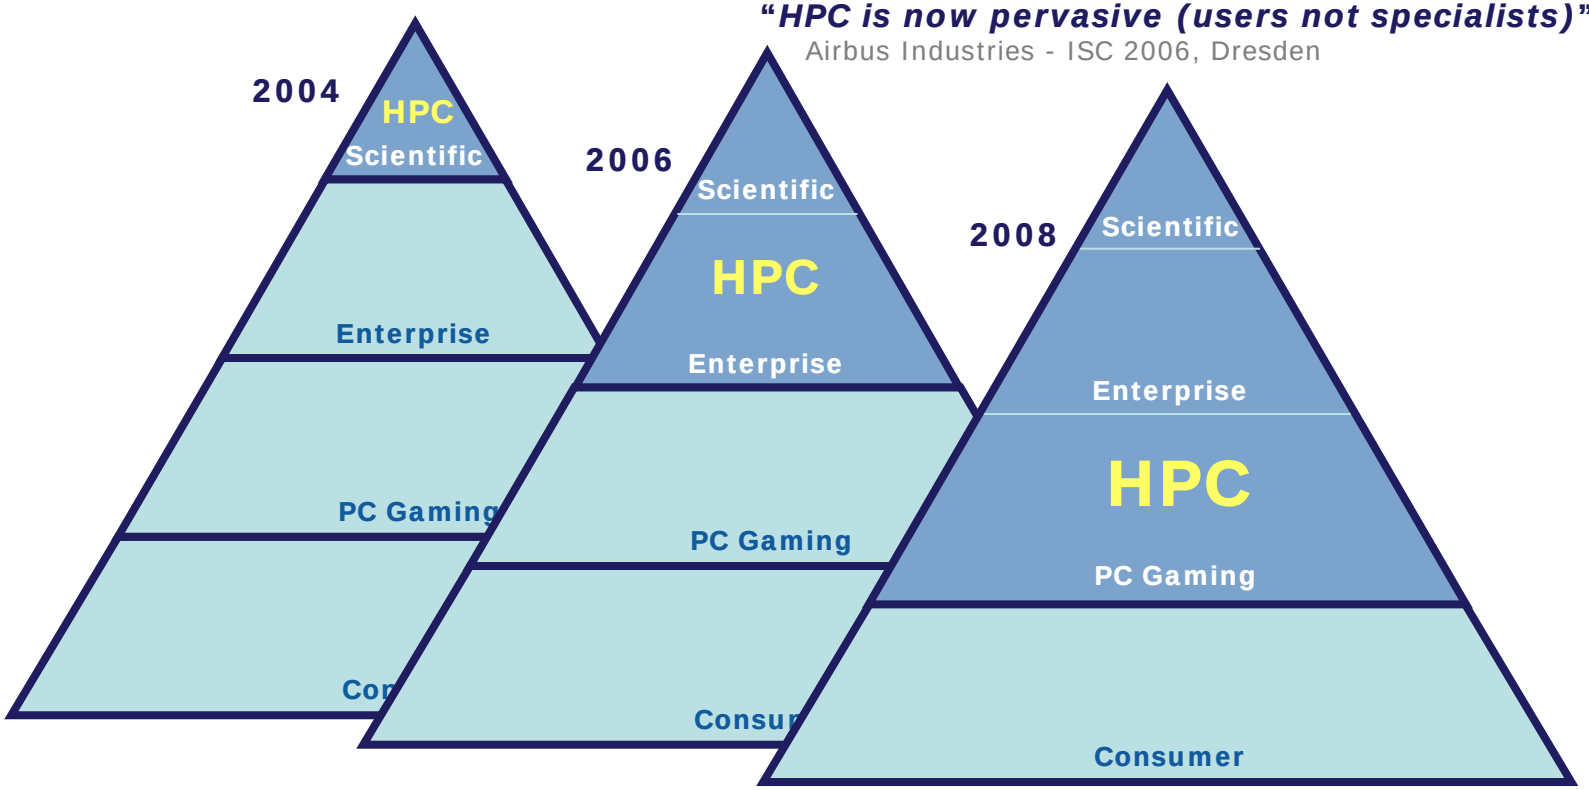

Commodity Technology Drives HPC Expansion Minimizes End-User TCO

Strong Demand for Commodity-Based HPC Infrastructure

HPC Server Revenue by Segments

Source: IDC - 2008

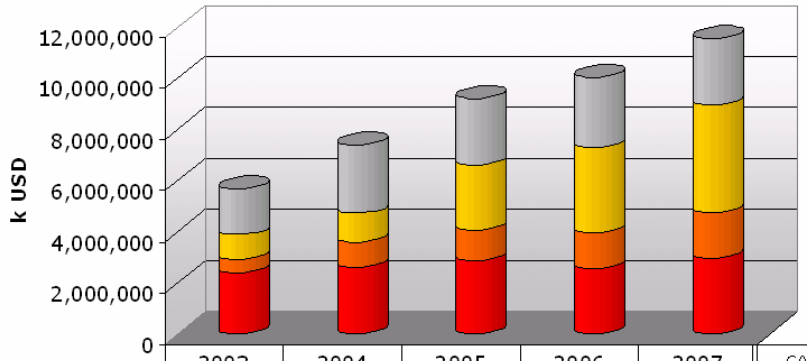

Price Range		2003	2004	2005	2006	2007	CAGR
<$99k	Workgroup	1,805,662	2,668,372	2,568,446	2,744,283	2,606,812	10%
$100k-$249k	Departmental	946,886	1,116,730	2,561,422	3,323,051	4,193,089	45%
$250-$499	Divisional	543,830	977,380	1,197,234	1,420,357	1,780,539	35%
>$500k	Supercomputers	2,401,478	2,630,877	2,881,105	2,566,886	2,982,811	6%

HyperTransport™ Carries Highest Values in HPC

by virtue of its

Low Latency, High Bandwidth, Direct Connect Architecture

Server Farms

Data Centers

Supercomputers

Gaming Platforms

HyperTransport™ CONSORTIUM

Performance and Scalability are Make-or-Break Factors

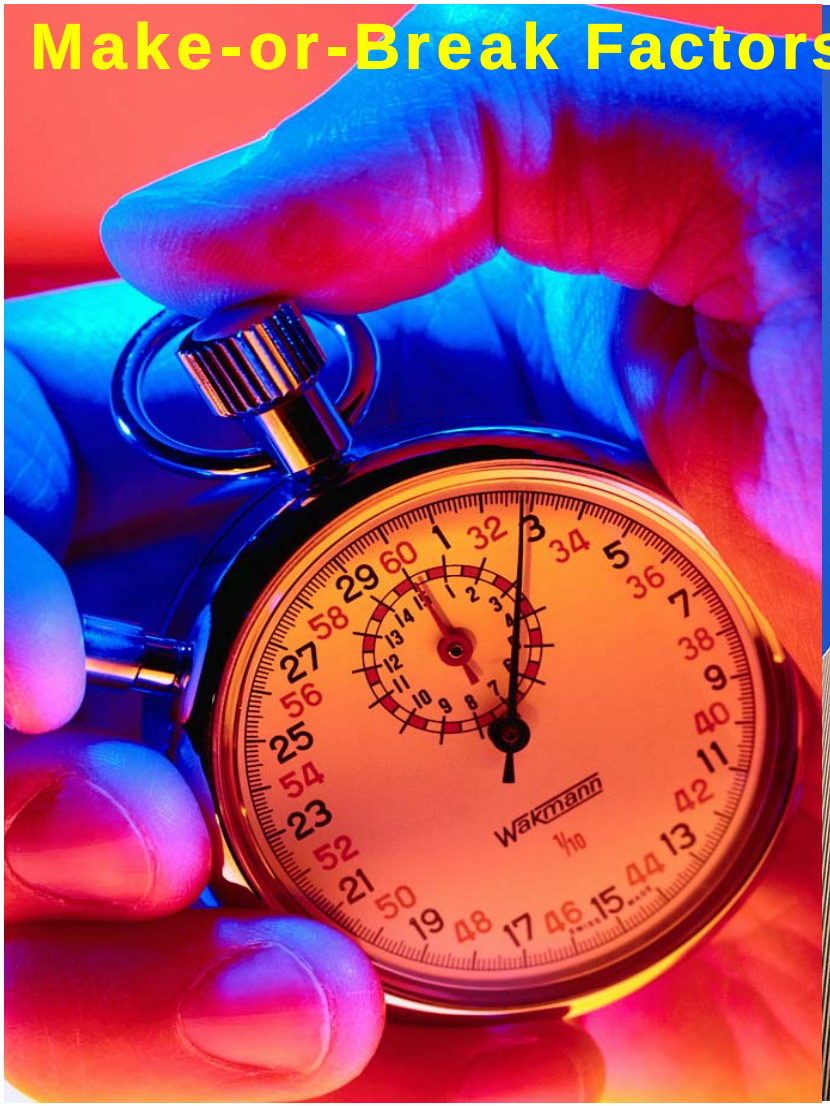

Interconnect Technology is Crucial to HPC Product Success

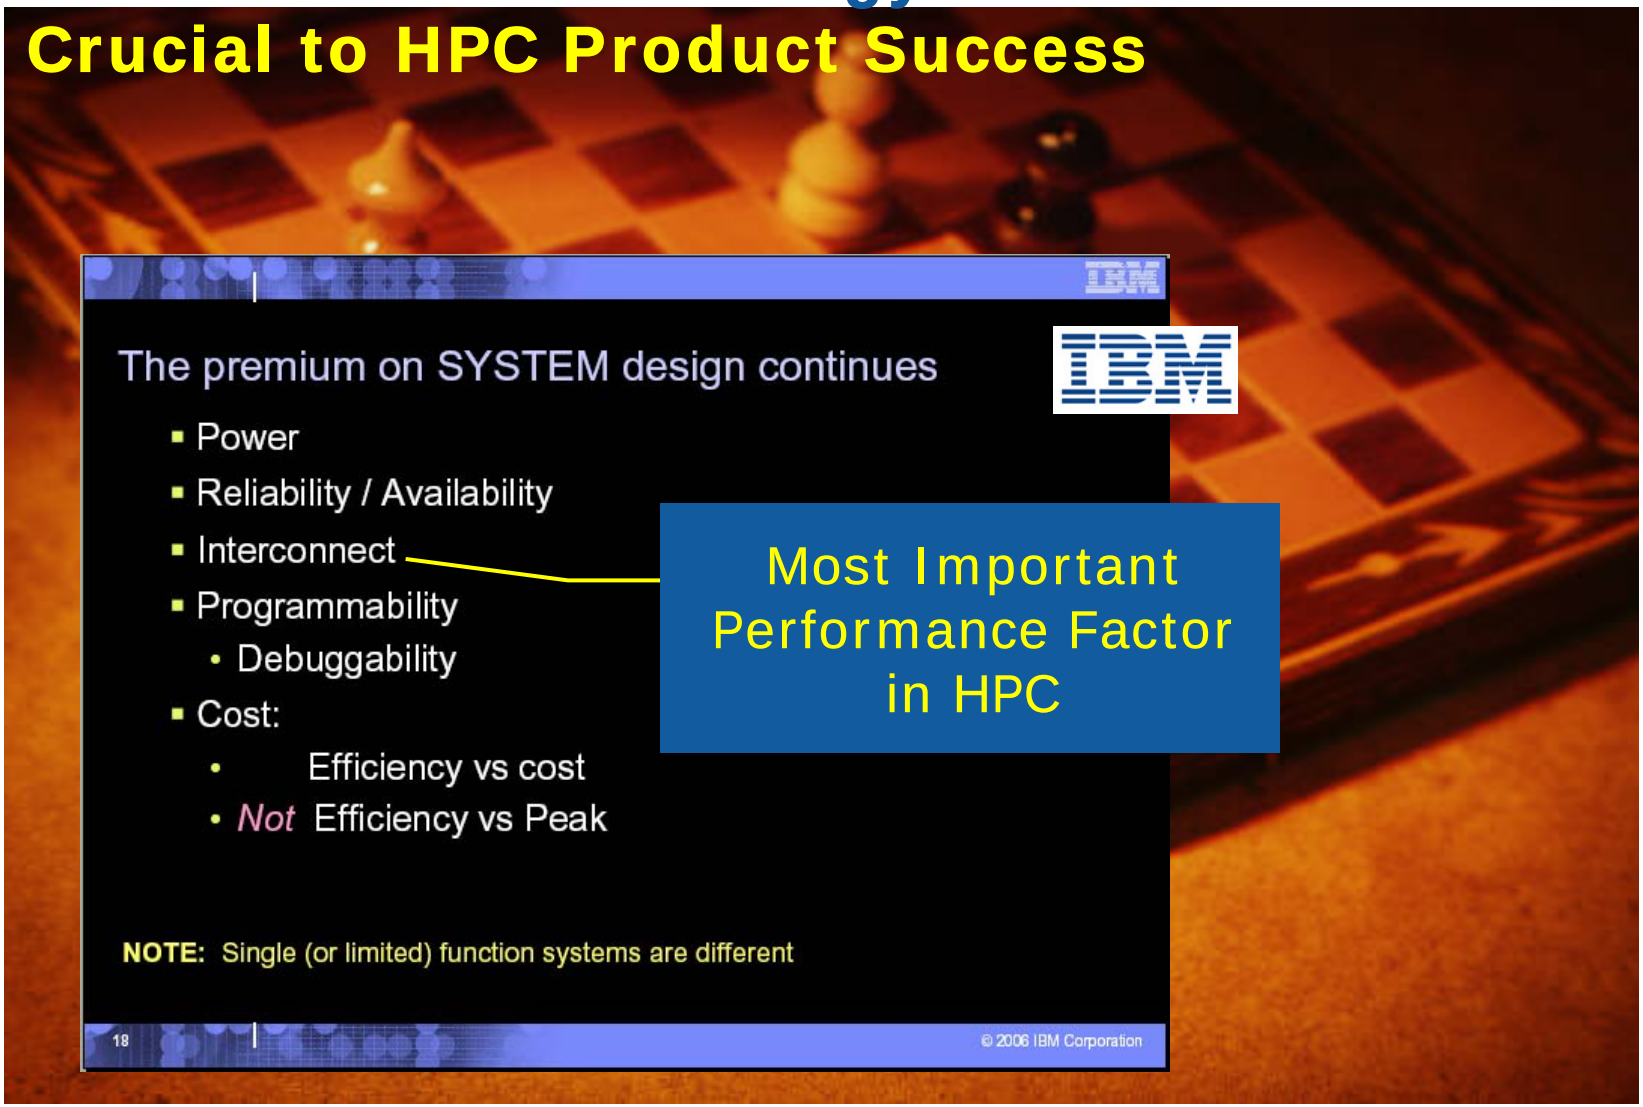

Interconnect Technology is Crucial to HPC Product Success (cont.)

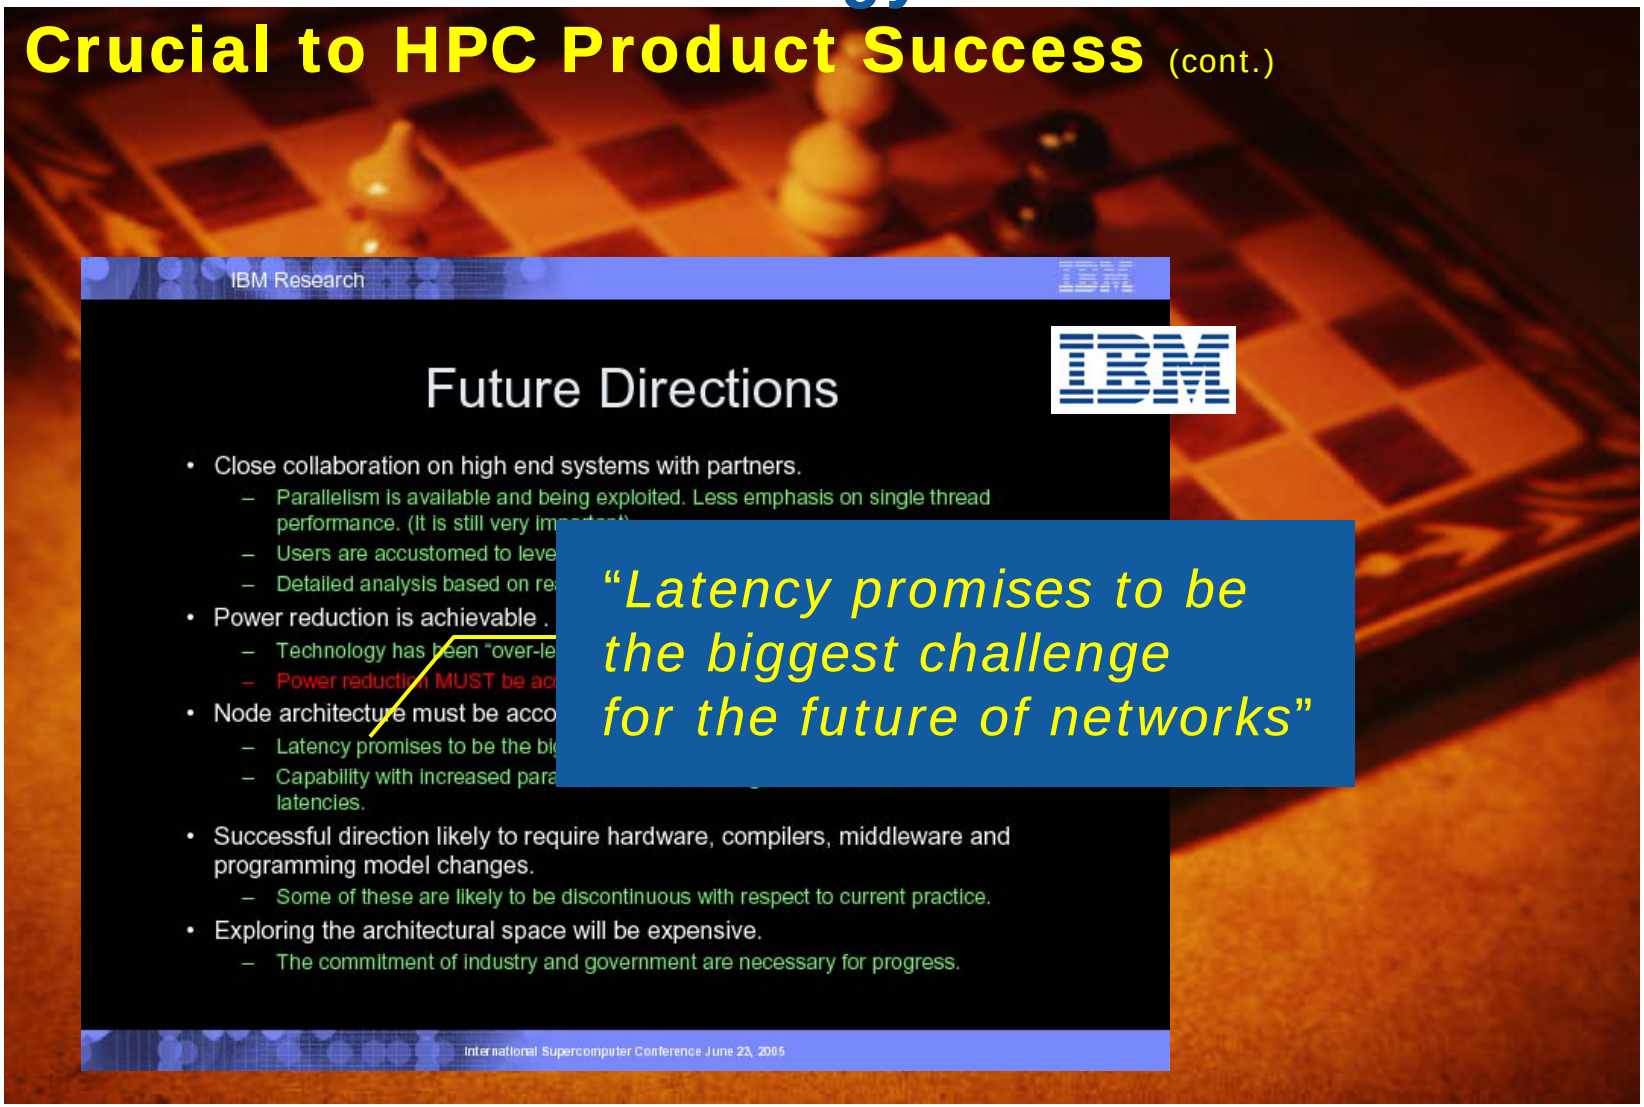

Latency is The HPC Market Driver

Peak computing power is no longer the metric to consider

Feeding processing units with data is the real concern, and flow control dependencies in applications make latency the key performance metric

Prof. Jose Duato
Polytechnic Univ. of Valencia, Spain (*)

HTX Has Been and Continues to Be The Lowest Latency Slot Interface in Industry

(*) Jose Duato is Professor of Computer Science at the Polytechnic University of Valencia, Spain and a world authority in high performance computing (HPC). Prof. Duato has directly contributed to some of the computer industry's most advanced HPC product developments by companies like Compaq, Cray, IBM and Sun Microsystems. In 2006 Prof. Duato received Spain's most prestigious technology achievement award from Queen Sophia of Spain. Dr. Duato is an active contributor to Hypertransport technology and the HyperTransport Consortium.

Technology Evolution

HyperTransport™ CONSORTIUM

HyperTransport™ Delivers Core Values to HPC Industry

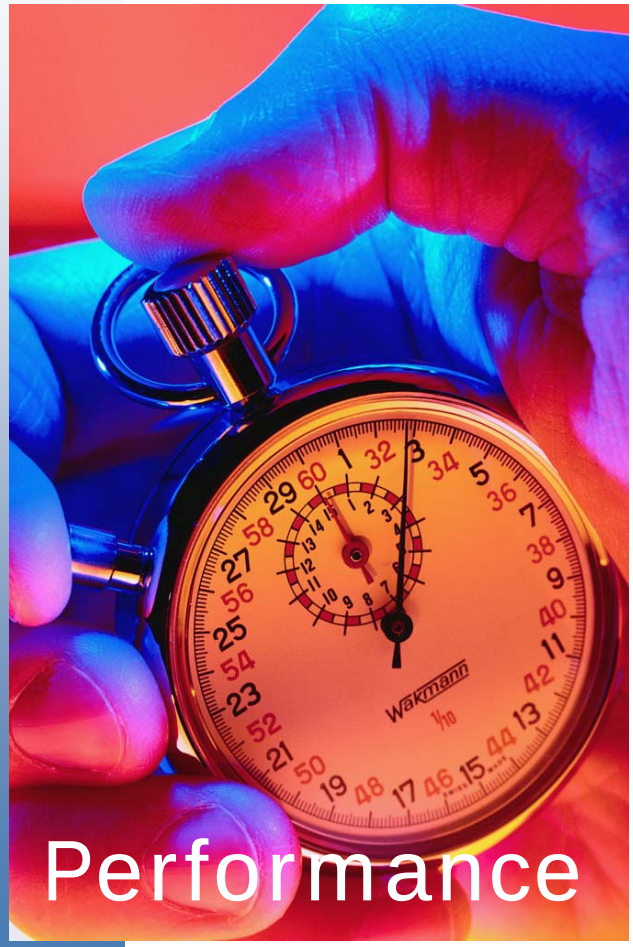

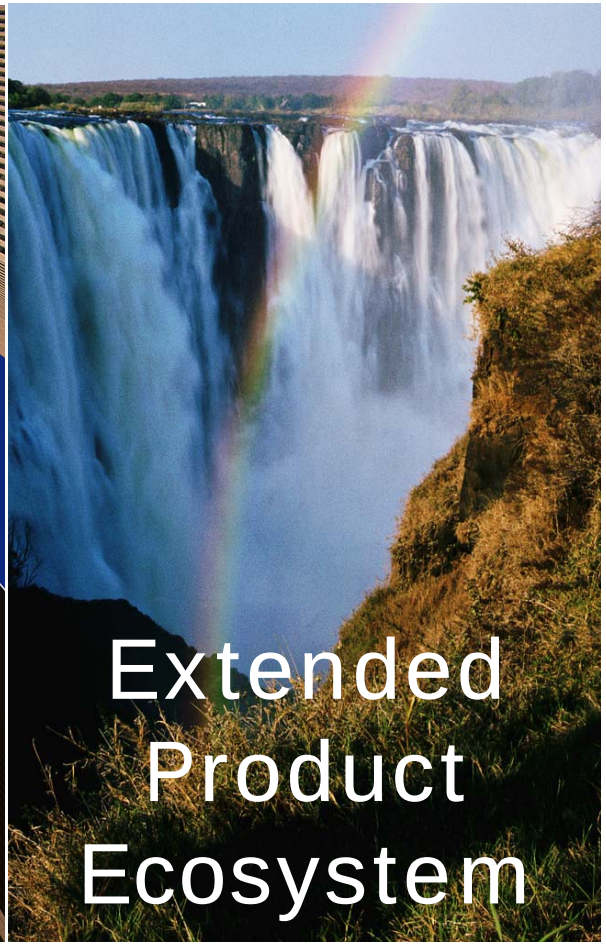

HyperTransport™ CONSORTIUM

Mature, Stable, Dependable Technology

HyperTransport™ is the Solution of Choice for Demanding Markets Sectors

58 Million HT-Powered Products

Capture	Market	Yr/Yr Growth
8%	Defense Applications	17%
32%	Top500 Supercomputers	28%
11%	Core Routers	1.2%
22%	Edge Routers	34%
15%	SAN	11%
23%	Servers	38%

Source: InStat - 2007

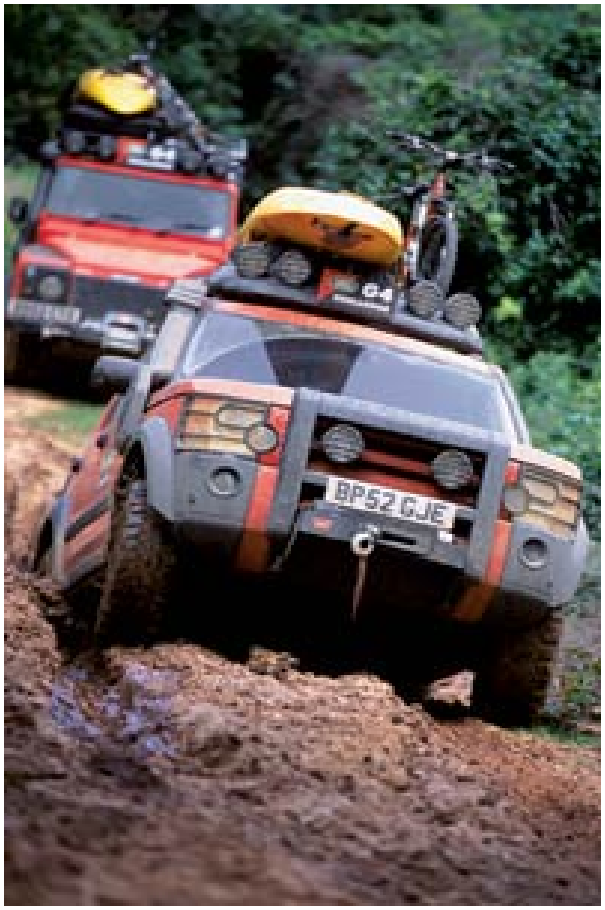

Mission-Critical Reliability

HTX™ Interface = Performance Edge

That PCI-Class Interconnects Cannot Deliver

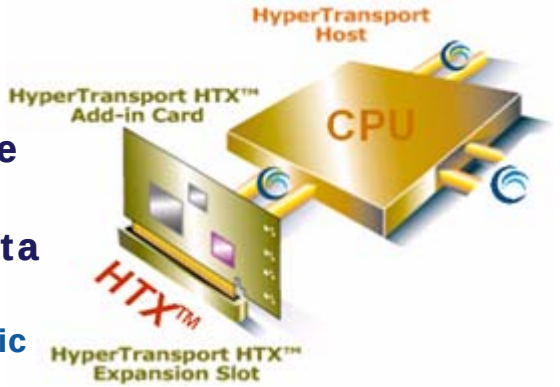

- **Lowest Latency, Highest Bandwidth Direct Connect Between CPU and Compute-Intensive Subsystems**
- **Removes Performance Bottlenecks in HPC Data Processing and Coprocessing Functions**
  - **No Intermediate CPU-to-Subsystem Control Logic**
- **Reduces Overall Power Consumption**
- **Complements PCI-Class Interconnects**
  - **Typically Co-Exist with PCI-X and PCI Express Connectors in Same MB Design**

Features

8-Bit and 16-Bit HyperTransport Link Support
800 MHz Max Clock Rate
6.4 GB/s Bandwidth (16-bit, Aggregate)

HyperTransport™ CONSORTIUM

HTX™ Leaps Forward to *HTX3*™

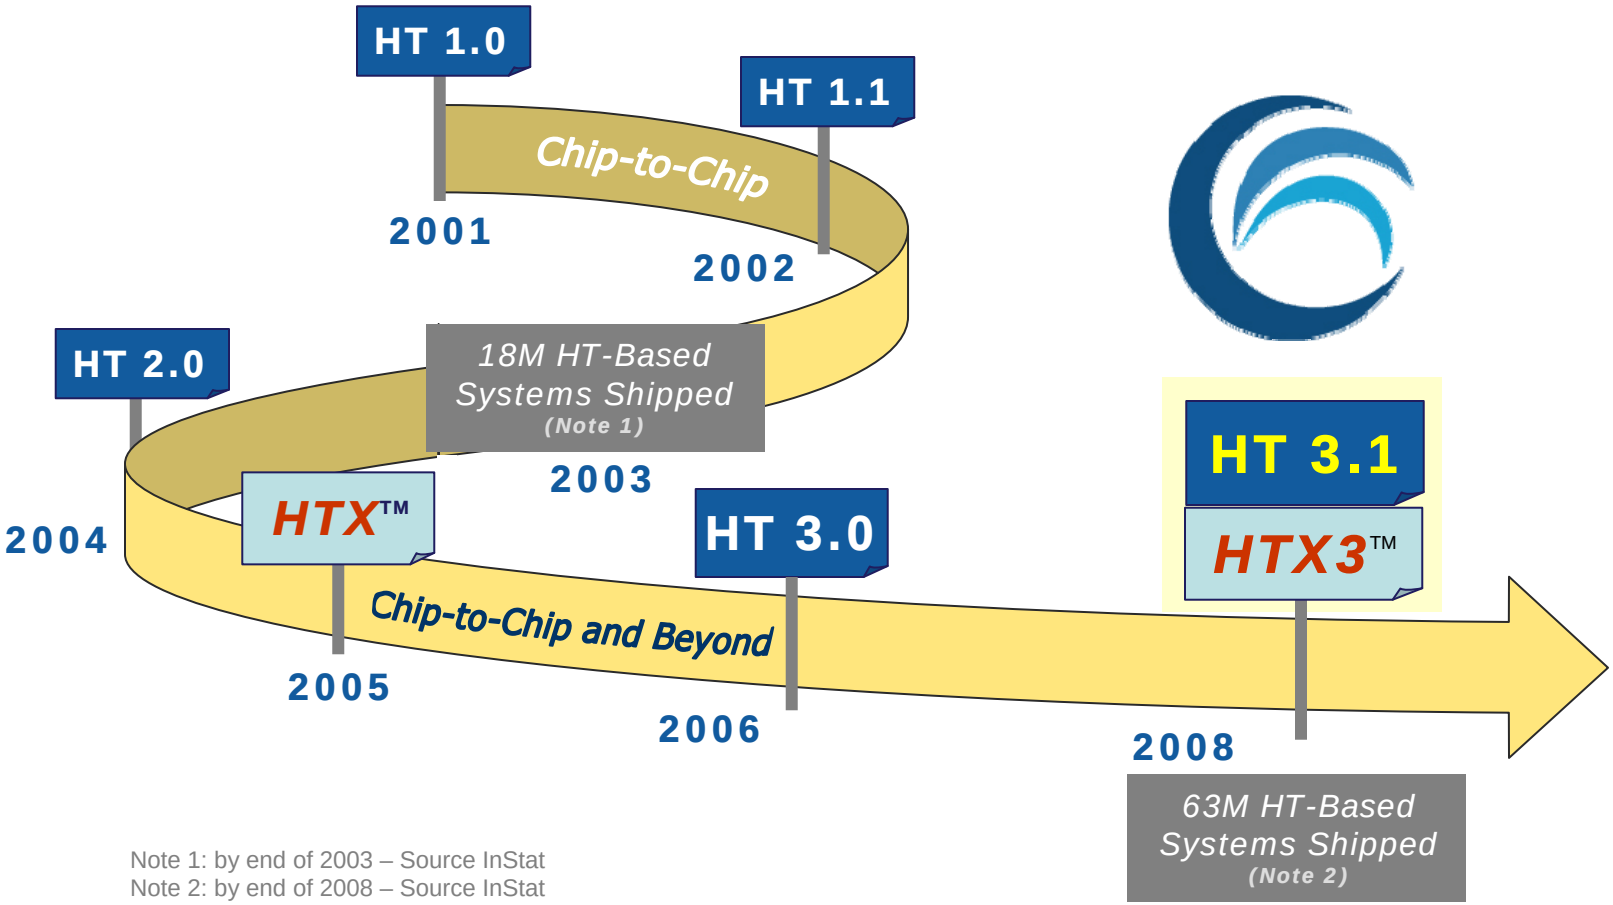

Note 1: by end of 2003 – Source InStat
Note 2: by end of 2008 – Source InStat

Introducing HyperTransport™ *HTX3™*

- >3x Bandwidth of Standard HTX Specification
- Brings HT 3.0 Performance and Architectural Latitude to HTX Connector
- Triples HT Clock Rate While Preserving Signal Integrity
  - Leverages HT 3.0 Clock Recovery “Scrambling” and “Training” Features
- Adds Power Management Features
- Adds Link-Splitting Capability
  - Allows Subsystem Connection to 2 CPUs

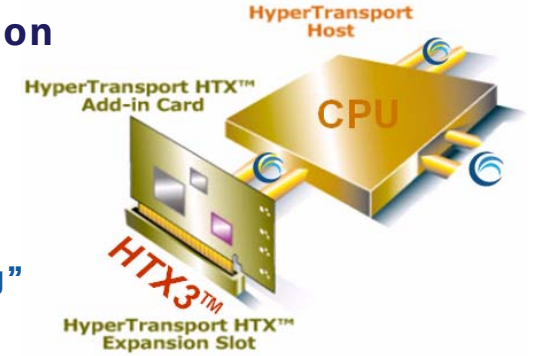

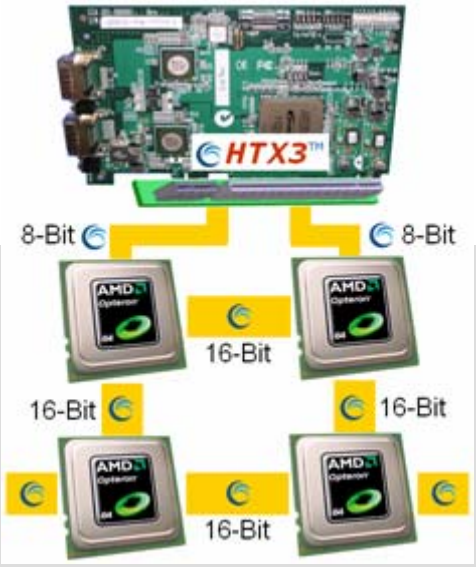

HyperTransport™ CONSORTIUM

HTX3™ Features Summary

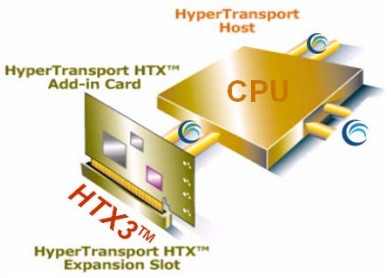

Feature	*HTX*	*HTX3*	Notes
Max Clock Rate	800 MHz	**2.6 GHz**	**12" Trace length**
Max Bandwidth x Lane	1.6 GT/s	**5.2 GT/s**	**Bi-directional**
Max Bandwidth Aggregate	6.4 GB/s	**20.8 GB/s**	**Bi-directional 16-Bit HT link**
HT3 Link Splitting Support	NO	**YES**	**HT link can be 1x 16-Bit or 2x 8-Bit for mutli-CPU support**
HT3 Extended Power Management	NO	**YES**	**LDTREQ# Signal Added to participate in x86 power states**
Extended FPGA Guidelines	NO	**YES**	**Incorporated field-proven recommendations**
Full Backward Compatibility	--	**YES**	**Level shifters and signal allocation**

Why HyperTransport™ *HTX3™*?

- FPGAs Play Key Role in Compute-Intensive Designs
- HTX3 Well Aligns with Expected FPGA Technology Progress
- FPGA Cores Expected to Support HT3.0-Class Clock Rates
  - From Bandwidth Bottlenecks to Drivers
- Power Optimization Ranks High in HPC Agenda
- HT 3.0 Has Reached Maturity and Stability
  - Design Learning Curve
  - Product Implementations
  - Market Feedback
- HT 3.0 Capability Now Safely and Stably "Connectorized"

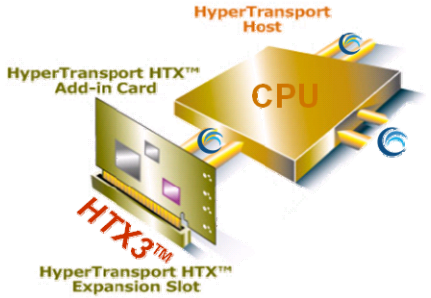

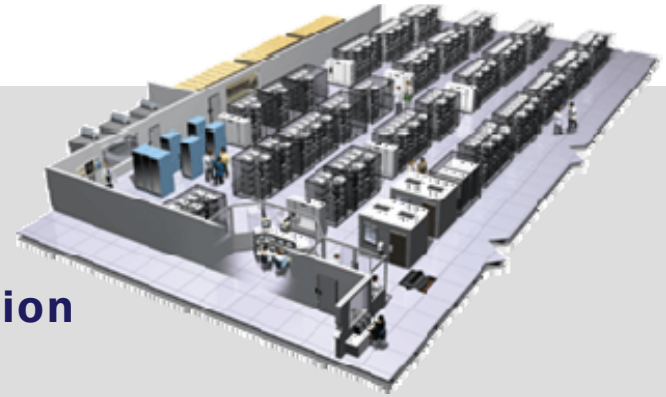

With its State-of the-Art Features Set *HTX3™* Empowers Future HPC Innovation

HyperTransport™ CONSORTIUM

Core Values

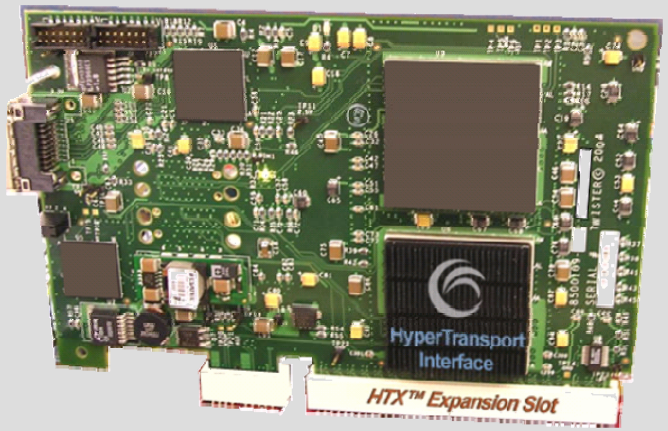

HTX™ and *HTX3*™ Connector

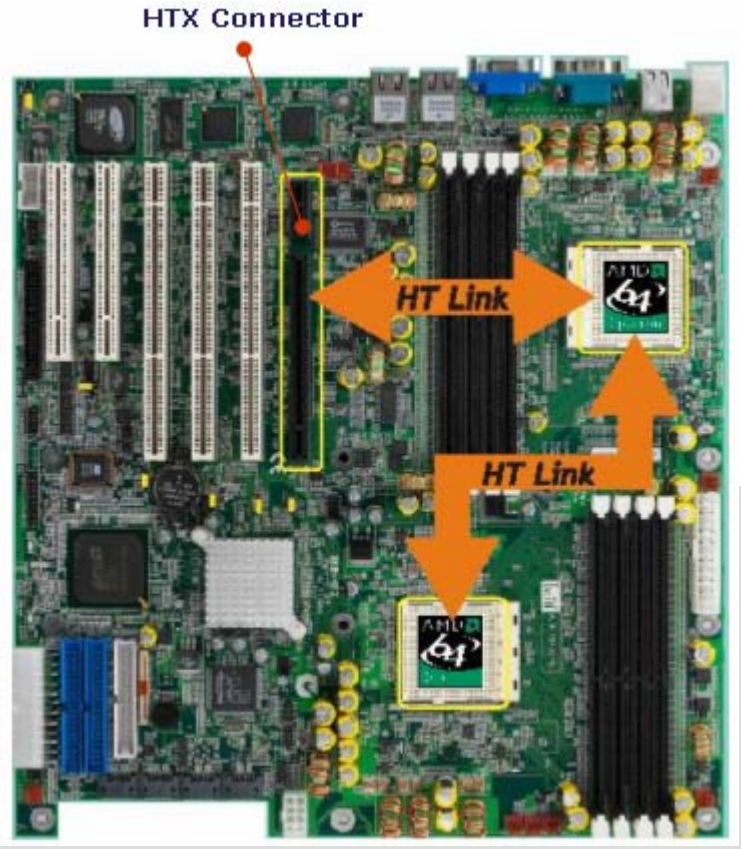

- **Uses popular PCIe Connector**
  - **High-volume part, Low cost**
- **HT-defined signals**
- **Reverse-Installed vs. PCIe Slots**
  - **Prevents Wrong Card Insertion**

HTX™ and *HTX3*™ Connector Mounting

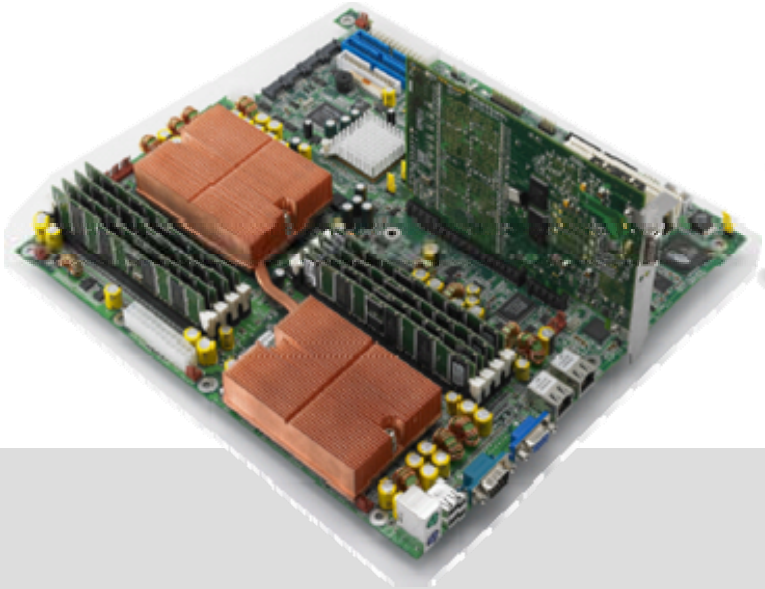

- **Vertical Mount**
- **Horizontal Mount via HTX Riser**

HyperTransport™ CONSORTIUM

Lowest Achievable Latency

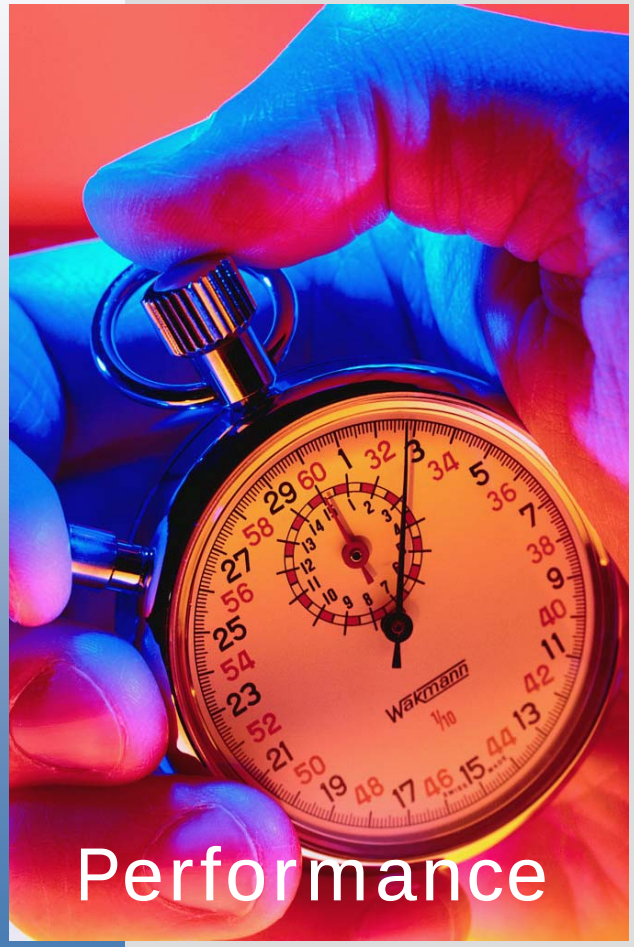

No Control Logic Penalty

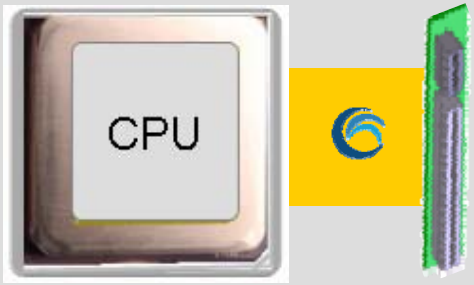

Control Logic Penalty

Lowest Achievable Latency (cont.)

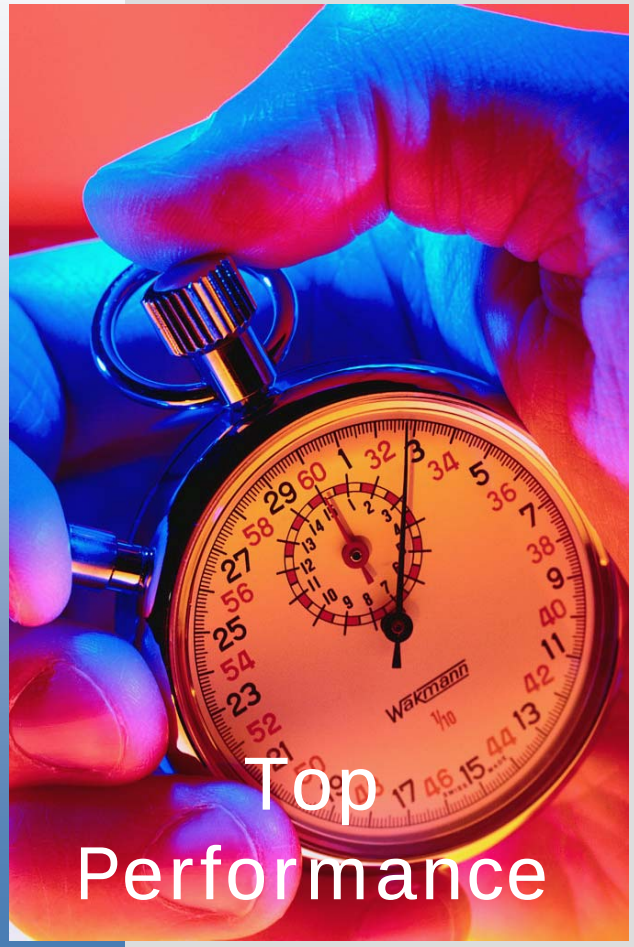

No 8B/10B Overhead Penalty

20% 8B/10B Fixed Overhead Penalty

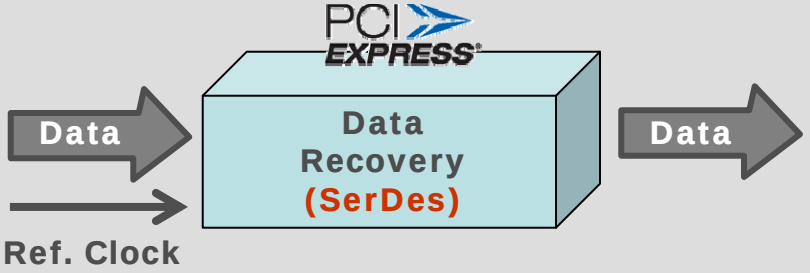

Lowest Achievable Latency (cont.)

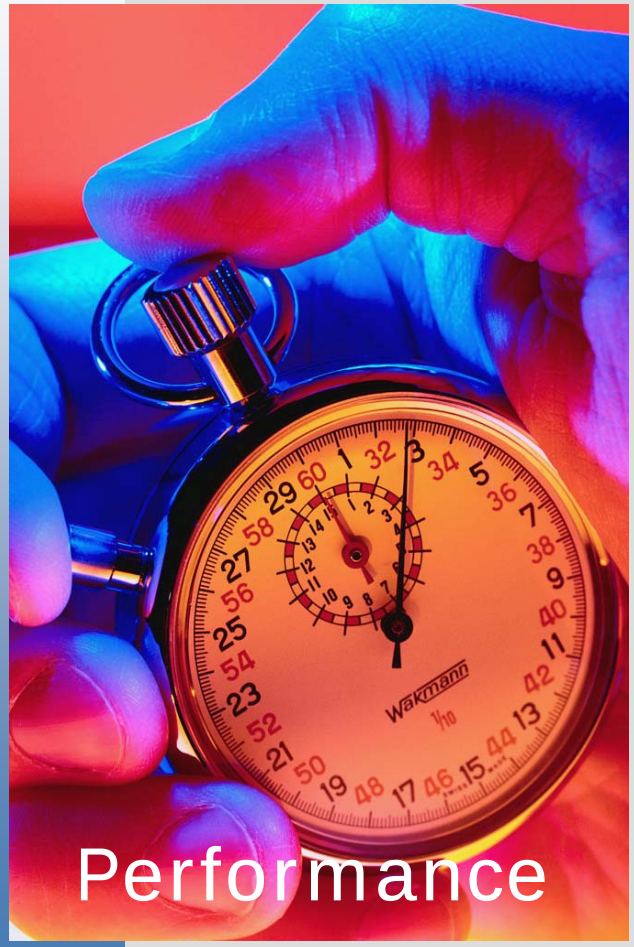

Minimized Packet Overhead

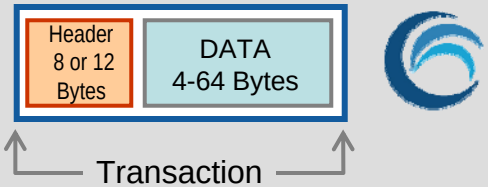

12 Byte Overhead Penalty vs. HT

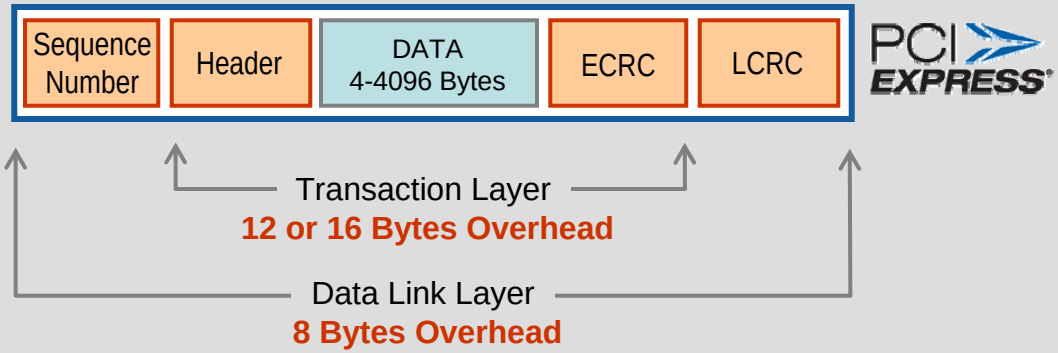

HyperTransport™ CONSORTIUM

HTX™ Delivers the Extra Performance that PCI-Class Interfaces Cannot Deliver

Complements PCIe and PCI-X

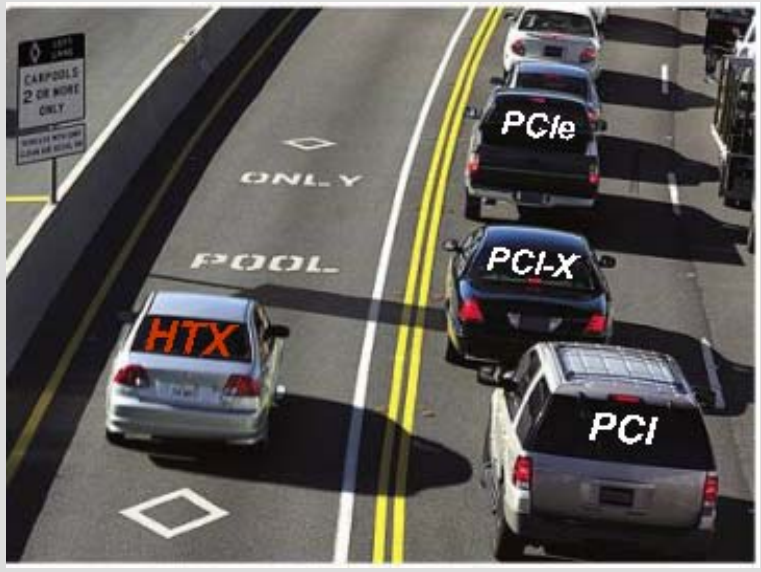

Cohexists within Same System

Main Applications

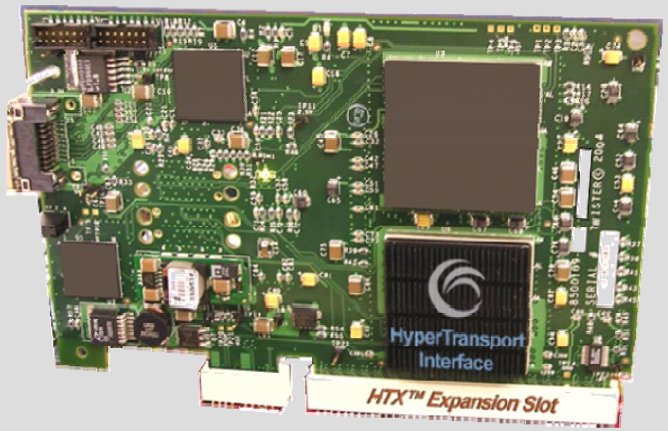

HyperTransport™ CONSORTIUM

HTX™ Applications

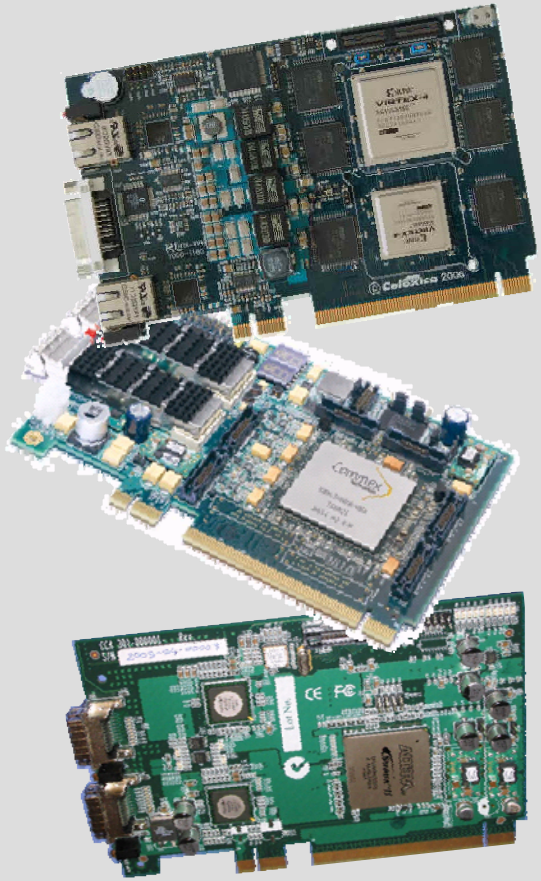

Compute-Intensive Board Level Subsystems Requiring:

- **High Bandwidth + Low Latency**
- **Multi-CPU Connections** (HTX3)
- **Advanced Power Management** (HTX3)

Targeting:

- **Database Analytics**
- **Stock Trading Acceleration**
- **High-Traffic Web-Based Applications**
- **Transaction Servers**
- **Streaming Media Servers**
- **Storage Servers**
- **Server Clustering and SMP**
- **Financial Modeling**
- **Communications**
- **Content Processing**
- **Security Processing**
- **High-Perf Real World Rendering**

HyperTransport™ CONSORTIUM

Product Ecosystem

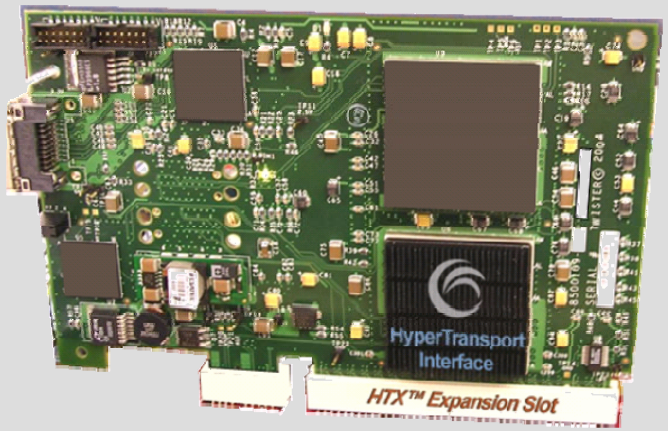

HyperTransport™ CONSORTIUM

Expanding *HTX*™ Products Ecosystem

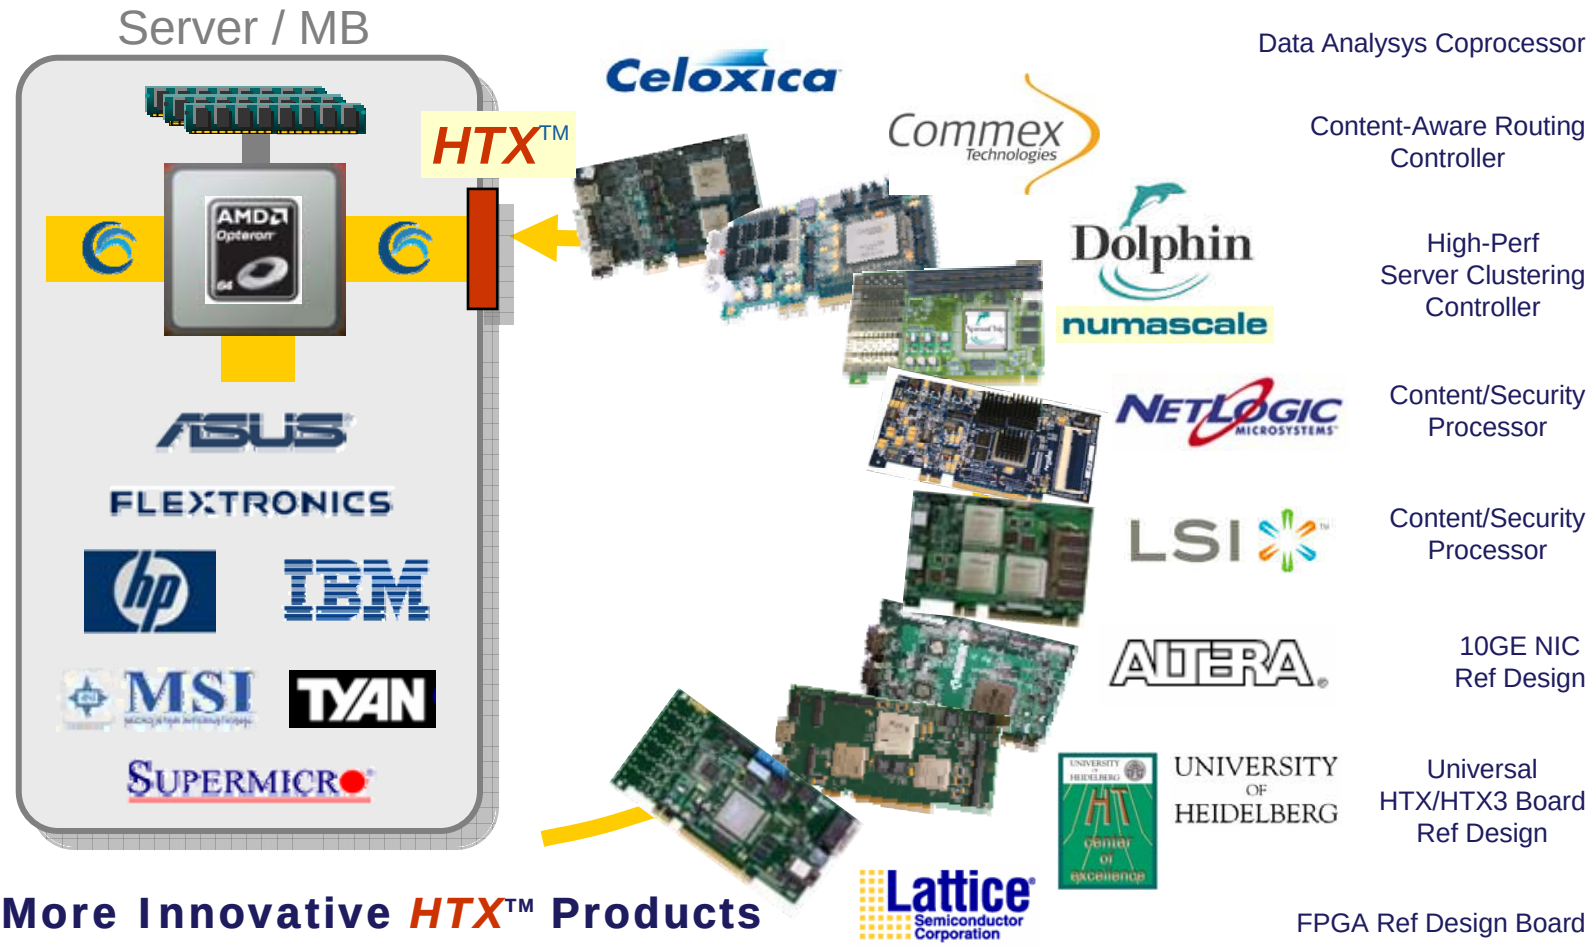

More Innovative *HTX*™ Products in the Pipeline

Tier-1 *HTX*™-Enabled Systems

ProLiant DL145-G3

ProLiant DL165-G5

ProLiant DL785-G5
2 HTX Connectors

New!

Tier-1 *HTX*™-Enabled Systems (cont.)

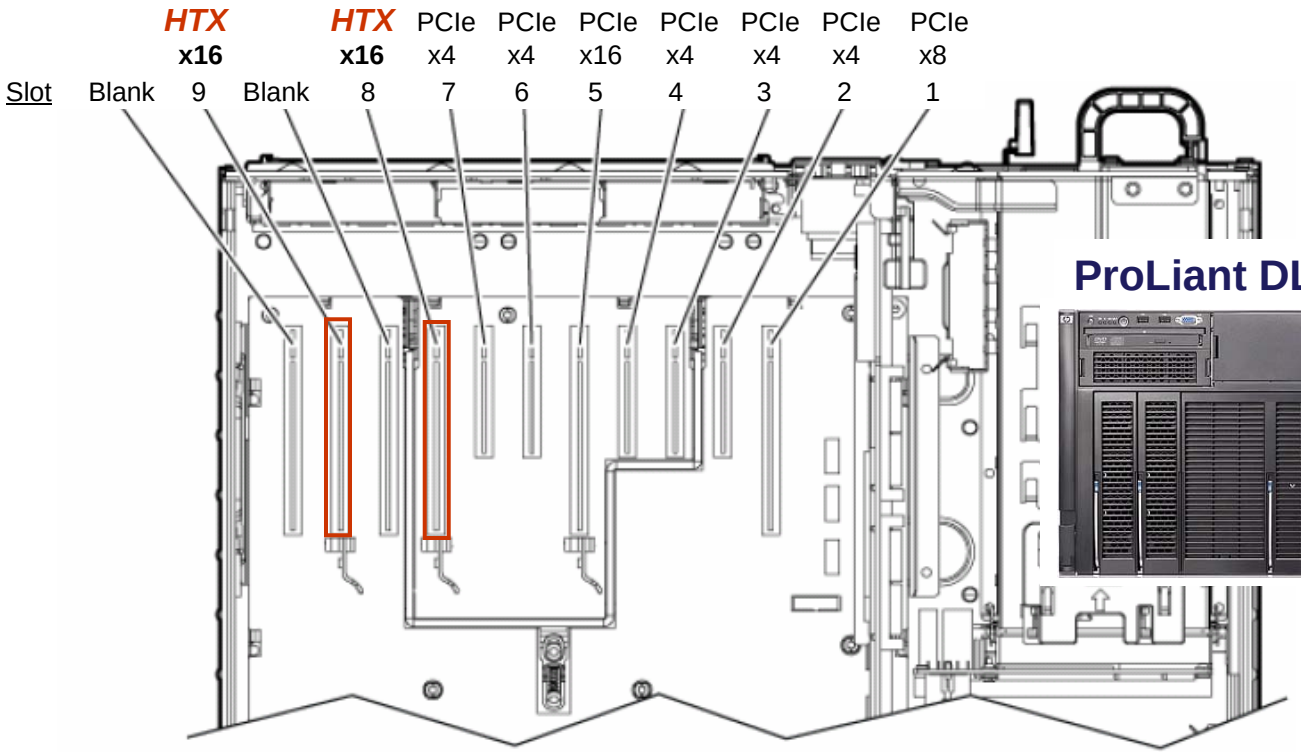

ProLiant DL785-G5

Tier-1 *HTX*™-Enabled Systems (cont.)

IBM

System x3455

System x3655

System x3755

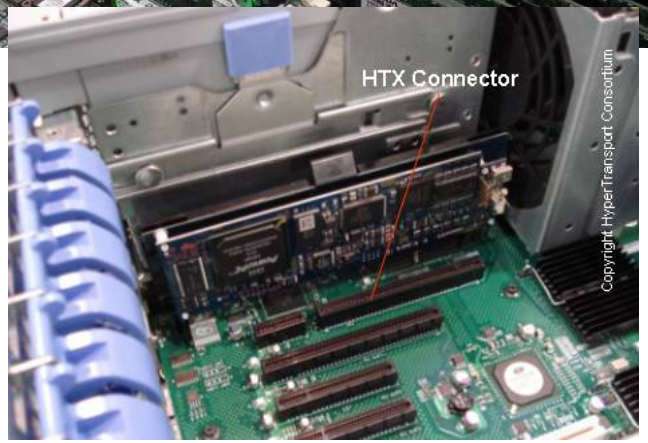

Tier-2 *HTX*™-Enabled Systems

SUPERMICRO

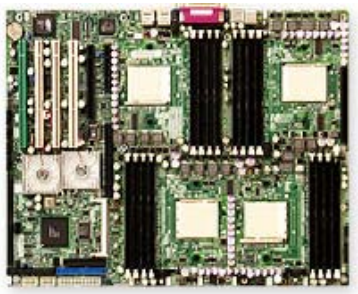

H8QC8
H8QC8+
H8QCE
H8QCE+

H8DCE

H8DSL

AS104C-8R

AS104C-8

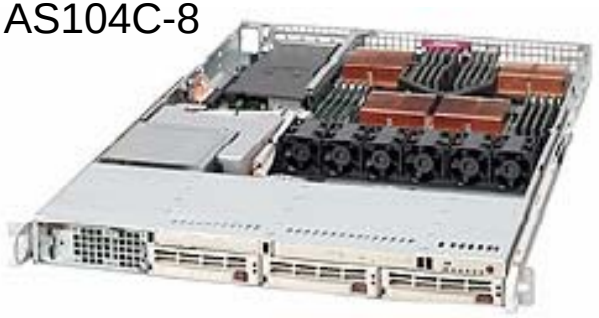

HTX Used Conventionally and for 4-way to 8-way CPU Expandability

HTX™ Subsystems

Celoxica — Market Data Capture, Life Sciences Accelerator

Packet Capture Technology

"Celoxica focuses on addressing some of the most challenging problems facing the finance sector, such as the exponential development of electronic trading and growth of complex derivatives trading, where the company can deliver advantage measured in microseconds."

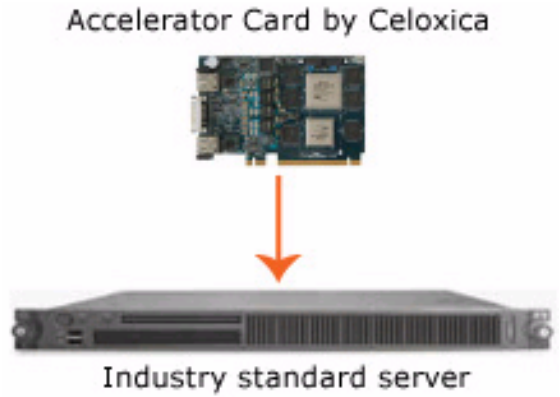

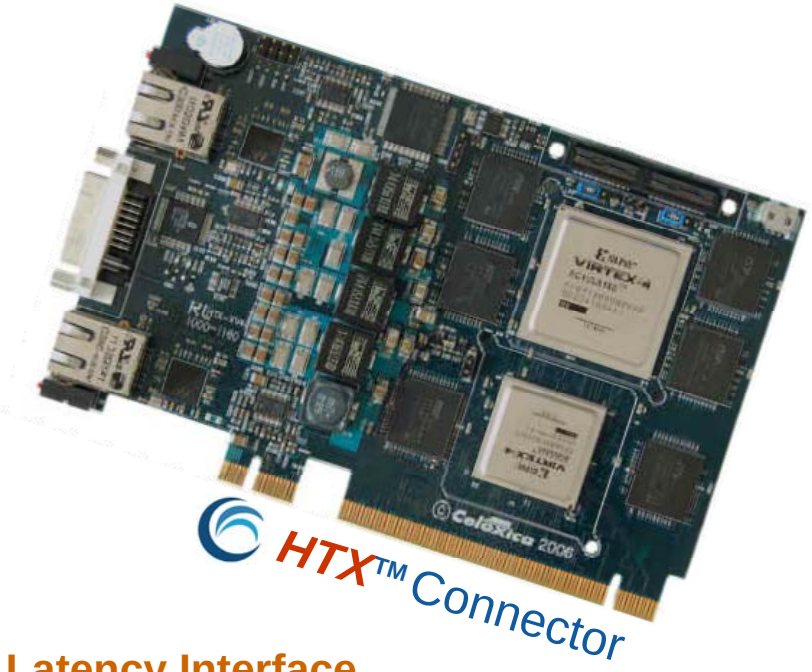

Lowest Latency Interface

"The HyperTransport [*HTX*™] bus is the lowest-latency bus available for communications between peripheral devices and a server's CPU."

HTX™ Subsystems (cont.)

Celoxica Market Data Capture, Life Sciences Accelerator

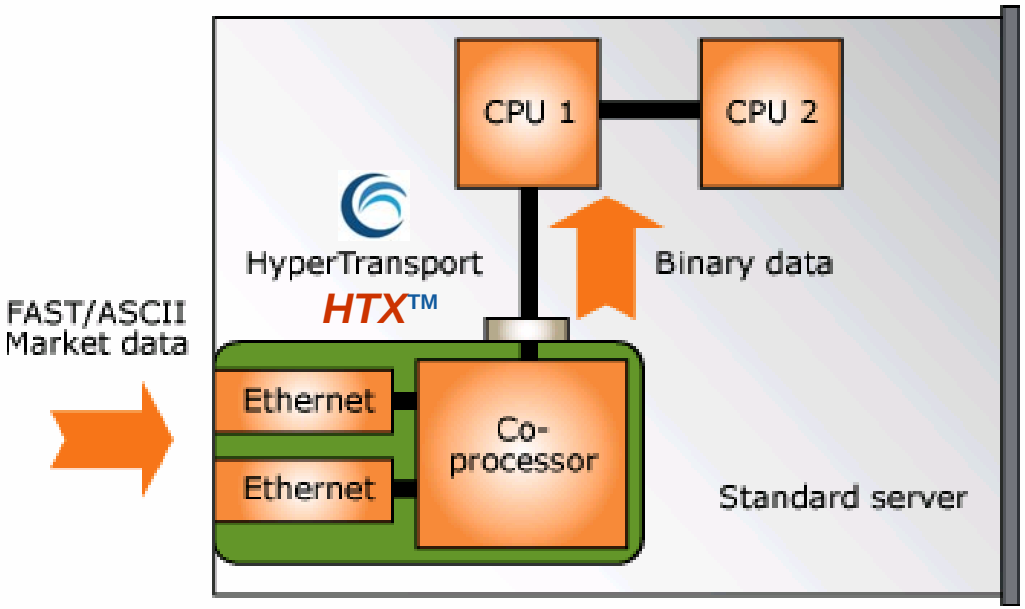

Innovative Architecture

"Celoxica's Accelerator is the first market data handler to directly connect a network interface to a co-processor, eliminating one of the major contributors to latency in a hardware/software co-processing system: the peripheral bus transactions between the co-processor and the network device."

Consistent Low Latency Regardless of Packet Size and Network Load

"Once converted to a binary representation, the network data is transferred via the HyperTransport bus directly to the CPU's memory. This transfer starts as soon as the data arrives off the wire, rather than buffering a complete packet before starting the transfer."

HTX™ Subsystems (cont.)

Celoxica Market Data Capture, Life Sciences Accelerator

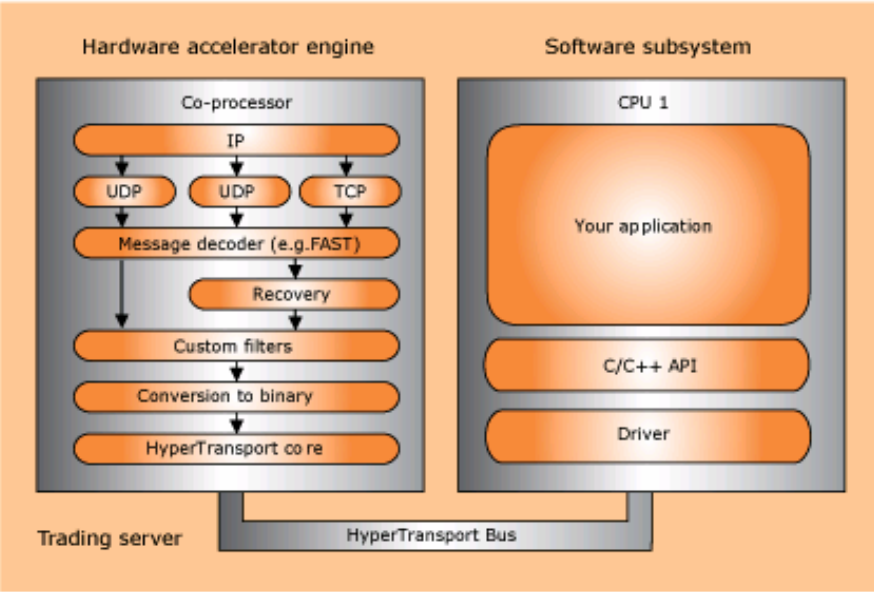

FPGA-Accelerated Protocol Handling

"The protocols that are handled in the FPGA include IP, UDP, TCP and the financial message format in use, which may be FAST, FIX or a proprietary format."

"Celoxica leverages the latest FPGA technology for the co-processing. The structure of FPGAs allows the sequential tasks of parsing several protocol layers to be executed in parallel, in what is known as a pipeline. A separate pipeline is run for each network port. This means that a complex protocol stack can reliably run at wire-speed without missing a single packet."

Celoxica Market Data Capture, Life Sciences Accelerator

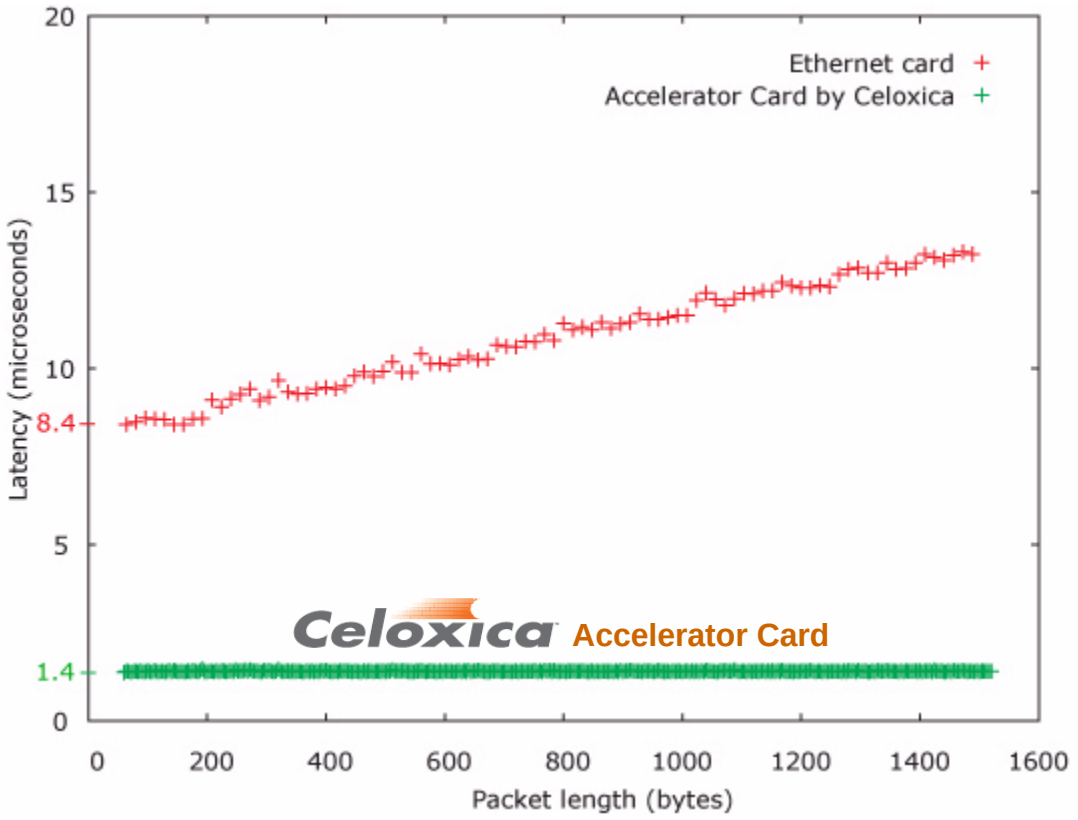

Constant Ultra-Low Latency

"The performance-optimized Ethernet card demonstrates a linear growth in latency with the size of packets.

The Celoxica Accelerator Card delivers constant ultra-low-latency access to network data regardless of packet size."

HTX™ Subsystems (cont.)

Celoxica Market Data Capture, Life Sciences Accelerator

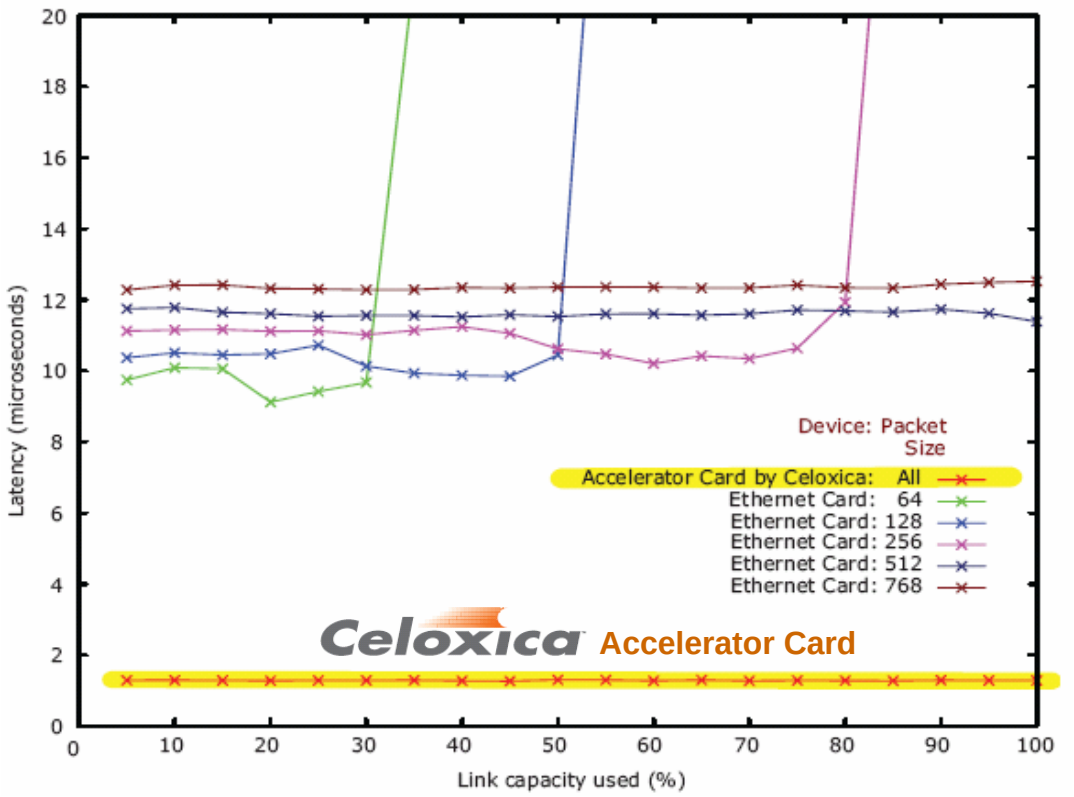

Load-Independent Latency

"The Ethernet card manage to give consistent latency for larger packets, but cannot keep up with bursts of smaller packets.The steep lines for 64, 128 and 256 bytes packets show that the latency increases rapidly as the system must buffer packets. The system rapidly starts dropping packets in these cases.

The Celoxica Accelerator Card is an important component for traders who need the lowest-latency access to market data, even during spikes in message volumes."

HTX™ Subsystems

Vulcan

Dual Port 10GbE Routing Processor for Multicore Systems

Commex Vulcan HTX Dual-Port 10 Gigabit Ethernet Server Adapter answers the ever-growing demand for application performance scalability, as it provides the application with only the network traffic it needs, when needed and delivered directly to the core that handles it. The result is better utilization of CPU resources and lower power consumption

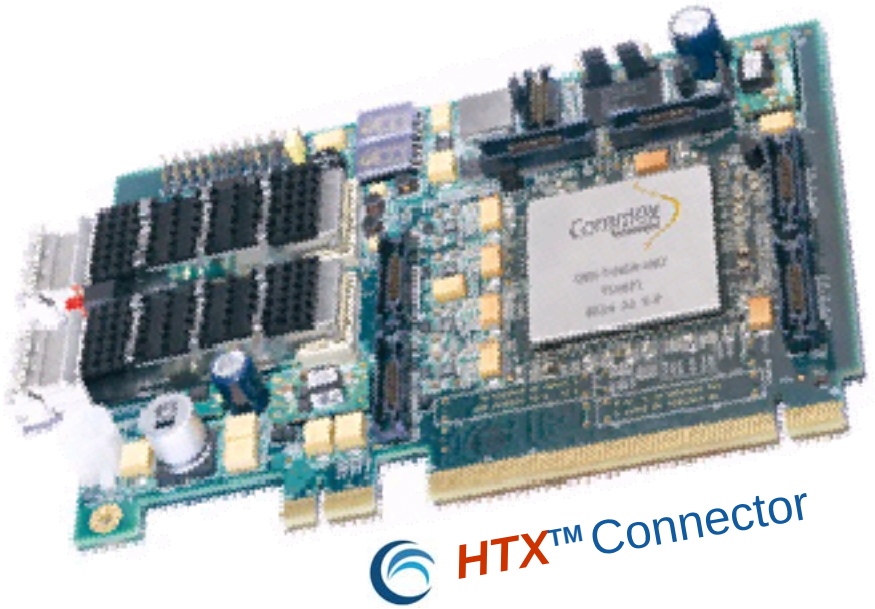

Applications

- **High-Traffic Web**
- **Telecom**
- **Automated Trading**
- **High-End AFEs & ADCs**
- **Data Center Applications Requiring High-Throughput, Fast Network Access**

HTX™ Subsystems (cont.)

Vulcan

Dual Port 10GbE Routing Processor for Multicore Systems

Equipped with Commex Thunder™ patent-pending hardware-implemented classification and action engine, Vulcan is capable of consistent association between CPU cores and their data streams, resulting in efficient multicore system load balancing of network traffic processing

Utilizing classification mechanisms such as pattern match (PM), Receive-Side Scaling (RSS), Smart Routing, multiple DMA channels, queues, and priority scheme, results in dramatically higher performance and reduced power consumption

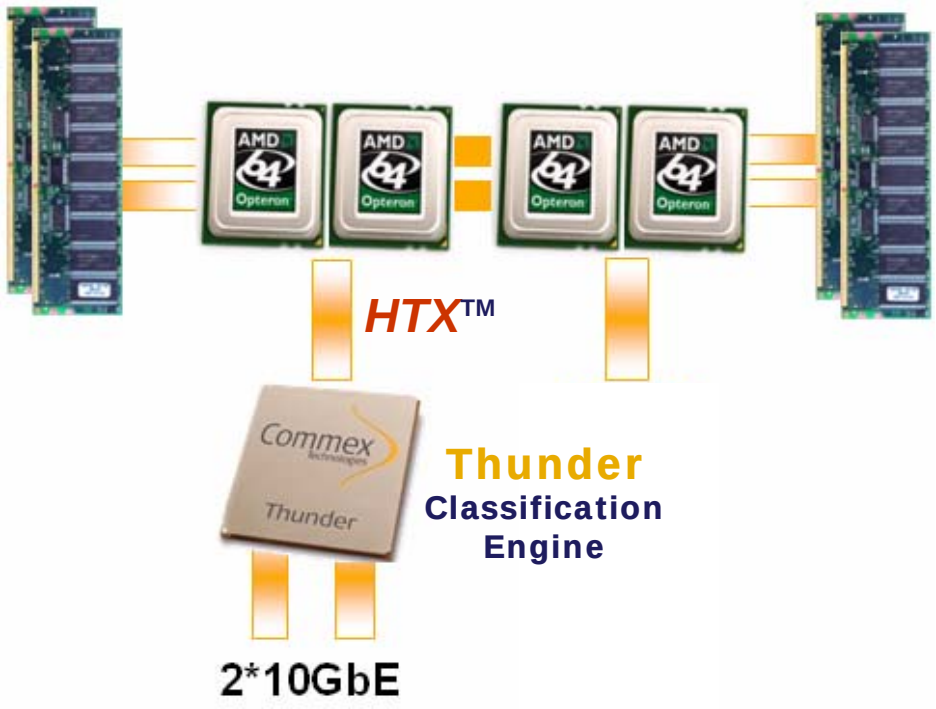

HTX™ Subsystems (cont.)

Commex Technologies

Vulcan

Dual Port 10GbE Routing Processor for Multicore Systems

- **Action**: Incoming network data classified based on format/pattern and forwarded to specific core using RX DMA channel
- **Benefit**: Network traffic processing performance scales with number of cores

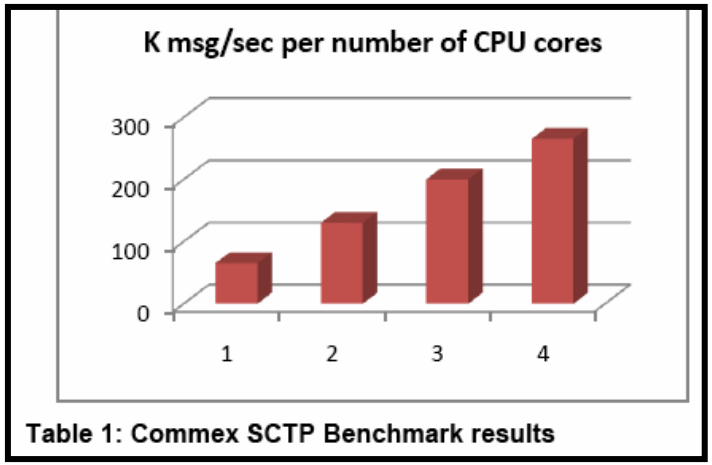

Table 1: Commex SCTP Benchmark results

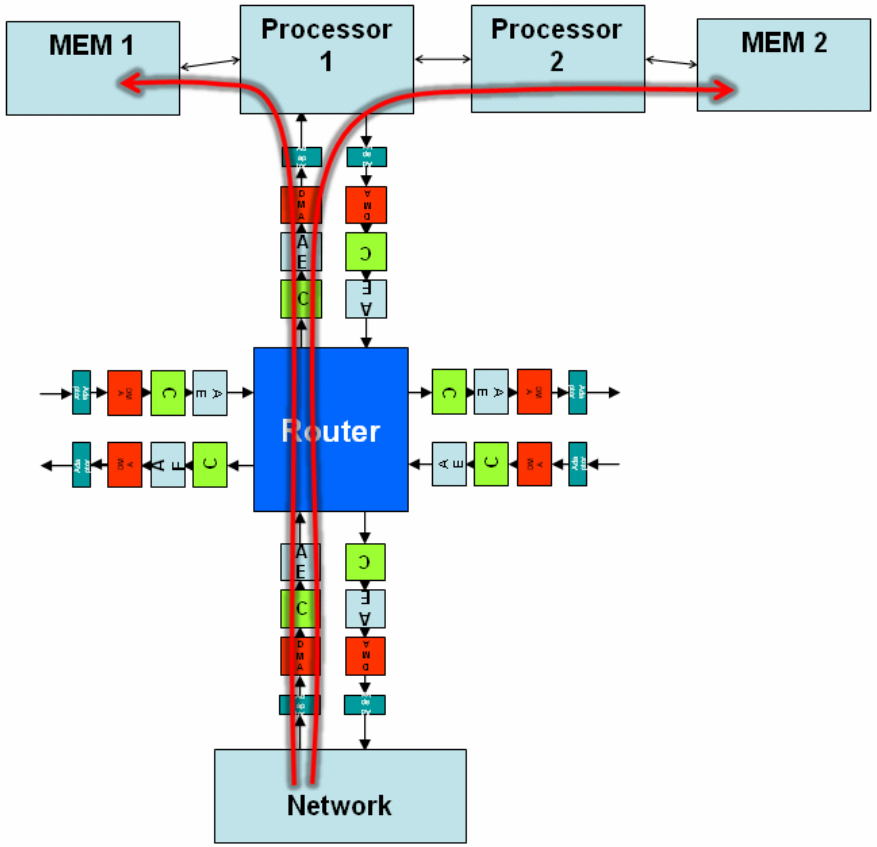

HTX™ Subsystems (cont.)

Commex Technologies

Vulcan

Dual Port 10GbE Routing Processor for Multicore Systems

- **Action**: Incoming network data classified based on priority and forwarded to specific core via QoS queue & DMA channel configured for proper priority
- **Benefit**: Multiple QoS schemes – e.g. simultaneous fixed low latency and high throughput

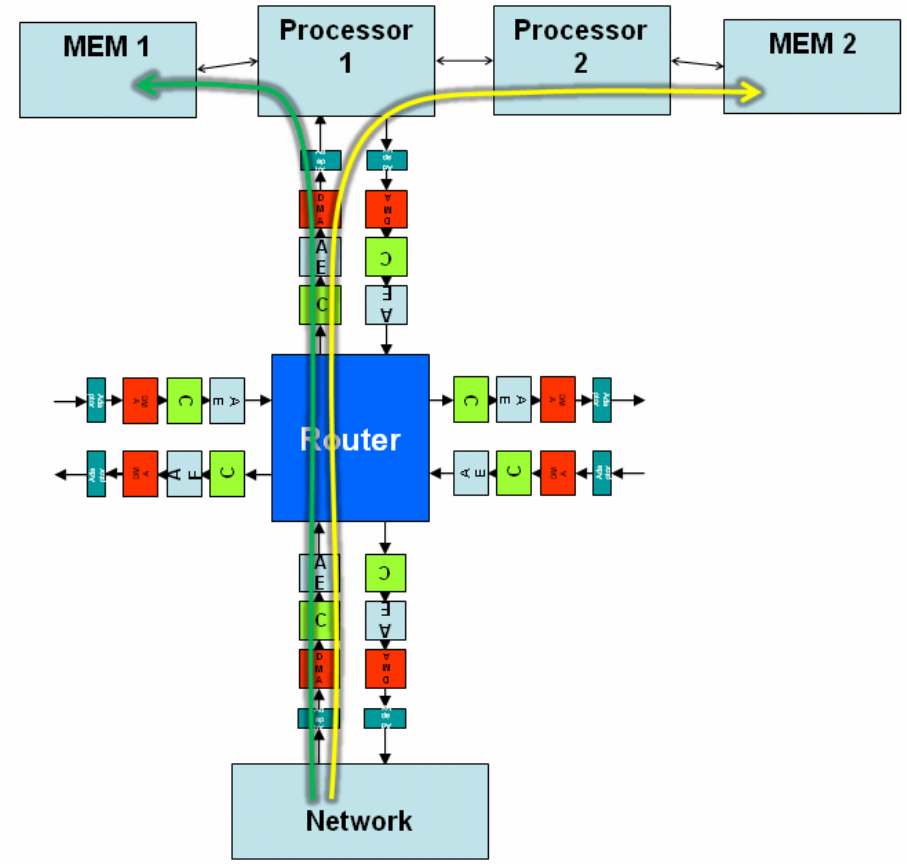

HTX™ Subsystems (cont.)

Commex Technologies

Vulcan

Dual Port 10GbE Routing Processor for Multicore Systems

- **Action**: Incoming traffic classified based on traffic type / platform load and selected traffic is re-routed back to network (content-aware bypass)

- **Benefits**:
- Bypass traffic overload to another server
- Bypass / drop irrelevant or "hostile" traffic
- Continue processing relevant service traffic within server throughput capability

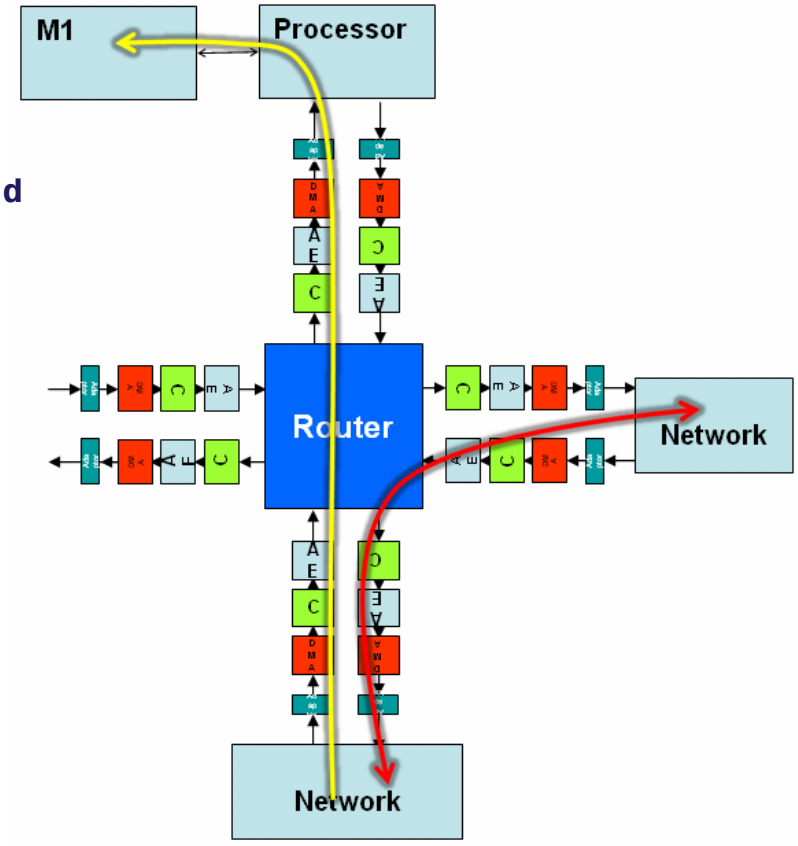

HTX™ Subsystems (cont.)

Vulcan

Dual Port 10GbE Routing Processor for Multicore Systems

Proven ability to continue streaming VoD under high request traffic VoD receives higher priority over new requests

Server A Leading Tier 1 Server	Server B Leading Tier 1 Server	Server C Leading Tier 1 Server with Thunder™
		
8 cores 2.6Ghz	4 cores 2.6Ghz	4 cores 2.6Ghz
1200 KPPS	850 KPPS	1500 KPPS
 0.54¢/PPS	 0.41¢/PPS	 0.24¢/PPS

Commex classification engine survives nearly double the request peak load for same server at nearly half the price

HTX™-Friendly Subsystems (cont.)

NLS1000 Content Processor

- Layer 7 Knowledge-Based Processor Family
- Optimized Inspection of Packets Traveling through Enterprise and Carrier-Class Networks
- Enables Real-Time Packet Processing Decisions
- 10 Gbps Line Rate Operation Across Multiple Packets

Specifications

Device	Description	Throughput	Interface(s)	Pacakage
NLS1000	NLS1000 Content Processor	10 Gbps	HyperTransport	1020 FCBGA
NLS1022	NLS1022 Payload Data Processor	10 Gbps	Proprietary	529 FPBGA

NETLT™
Knowledge-based Processors

Antivirus Gateways
Layer 7 Switches and Routers
Intrusion Detection/Prevention Systems
Unified Thread Management (UTM) Appliances

(*) HTX Proto/Evaluation Vehicle Available

HTX™-Friendly Subsystems (cont.)

NLS1000 Content Processor

NETLT
Knowledge-based Processors

- Development Kit for NLS1000 Layer 7 Content Processor
- HyperTransport HTX Interface
- Device Drivers
- High Speed Signature Compiler

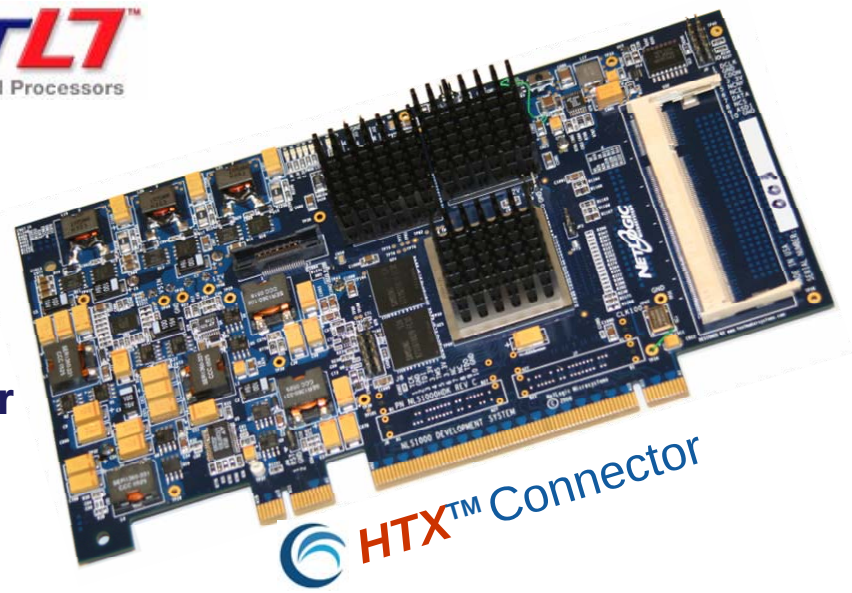

Layer 7 Content Processing Development Platform

HTX™-Friendly(*) Subsystems (cont.)

RMI CORPORATION

XLR Content Processor Families

- **Multiprocessing, Multithreaded Processing**
  - Throughput-Optimized MIPS64 Cores
  - 32 Fine-Grain Thread Processing
- **800 MHz, 1.0 GHz and 1.2 GHz Operation**
- **Leverages Macro Level Parallelism**
- **System and Network on Chip**
- **Autonomous Security Acceleration**
- **10 Gbps Cryptography**
- **General Purpose Programming**

Converged Networks and Computing Applications

(*) HTX Proto/Evaluation Vehicle Available

HTX™-Friendly Subsystems (cont.)

XLR Content Processor Families

- **Arizona ATX-II XLR Processor Evaluation Board**
- **Supports Complete XLR Processor Families**
- **ATX Form Factor**
- **Key XLR Processor Interfaces**
  - DDR1, DDR2 or RLDRAM Memory
  - TCAM/NSE
  - HyperTransport ***HTX***™ Connector
  - PCI-X Connector
  - Quad Gigabit Ethernet MACs
  - Dual 10Gbps Ethernet MACs
  - Firmware, OS and Debug Tools for Complete Test Environment

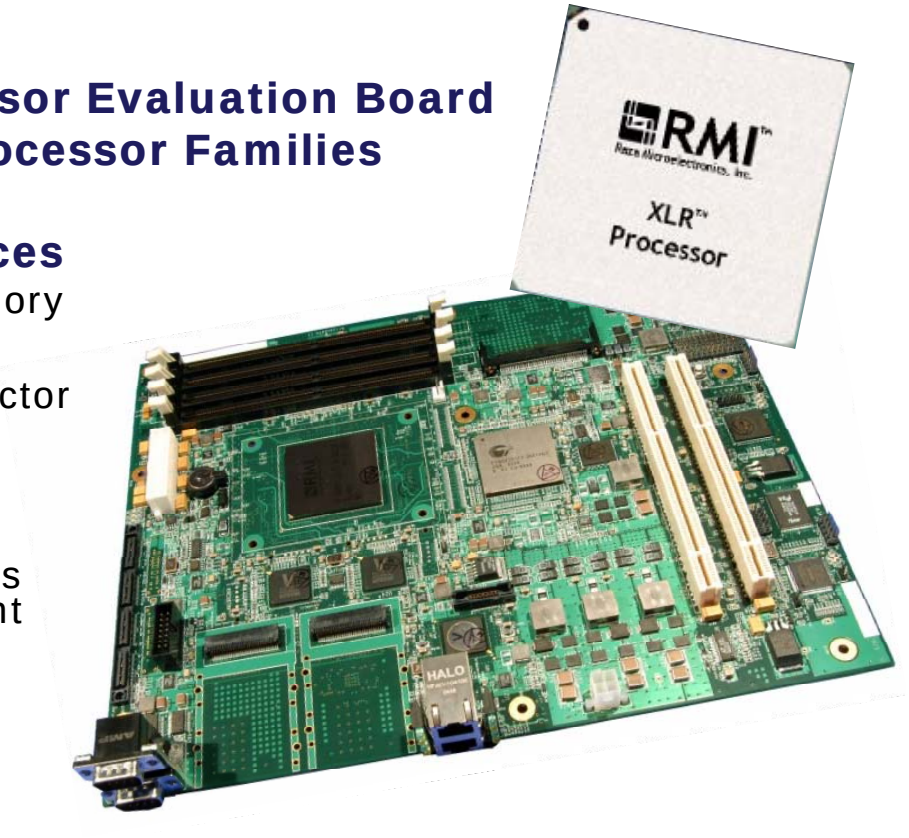

HTX™-Friendly(*) Subsystems (cont.)

Tarari T10 Content Processor

- **100 Mbps to 10 Gbps Real Time Performance**
  - Flow-Through Processing
  - Multi-Stream Capability
  - Efficient Multi-Threading
  - Cross-Packet Inspection
- **No Costly SRAM Required**

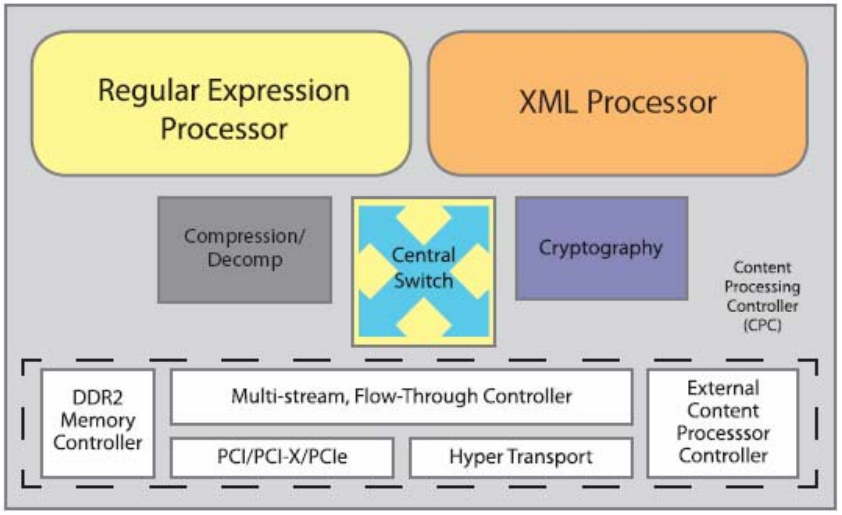

Powerful Data Analysis Network Security Engine

(*) HTX Proto/Evaluation Vehicle Available

HTX™-Friendly Subsystems (cont.)

Tarari T10 Content Processor

		T10 TECHNOLOGY FAMILY		
		T101xx	T102xx	T105xx
RegEx	Performance	2.5 Gbps	5 Gbps	10 Gbps
	Cross packet	Hardware	Hardware	Hardware
	# Complex Rules	1M	1M	1M
XML/ XTM	Performance	2.5 Gbps	5 Gbps	10 Gbps
	Cross Packet	Yes	Yes	Yes
Crypto	Performance	2.5 Gbps - AES	5 Gbps - AES	10 Gbps - AES
	Algorithms	AES, DES, 3DES, SHA1, RSA	AES, DES, 3DES, SHA1, RSA	AES, DES, 3DES, SHA1, RSA
Compression/ Decompression	Performance	2 Gbps	2 Gbps	2 Gbps
	Algorithms	GZIP, ZLIB, Deflate	GZIP, ZLIB, Deflate	GZIP, ZLIB, Deflate
	DMA	T10 Flow-Through Processing	T10 Flow-Through Processing	T10 Flow-Through Processing
	Technology	T10	T10	T10
	Bus	4-lane PCIe	- 16 bit Hyper-Transport - 8-lane PCIe	- 16 bit Hyper-Transport - 8-lane PCIe
	Power	3 W	4 W	5 W

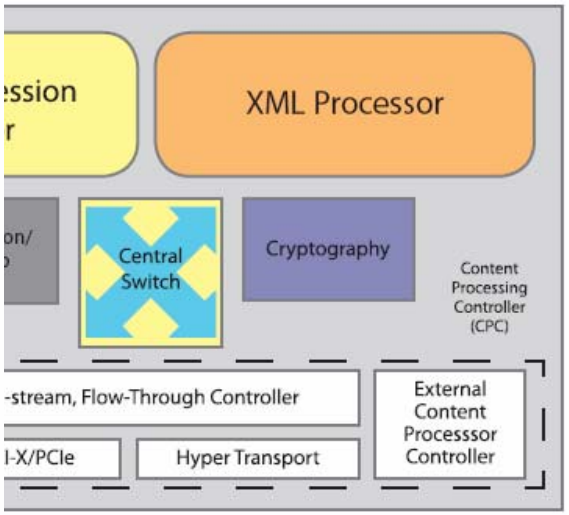

HTX™-Friendly Subsystems (cont.)

Tarari T10 Content Processor

- Tarari T10 Content Processor-Based
- 1-10 Gbps PSIX and PCRE Throughput
- 1 Trillion+ Matches per Second
- Advanced RegEx Features
- Advanced XTM/RAX/XML Features

High-Performance T10 Content Processing Development Kit

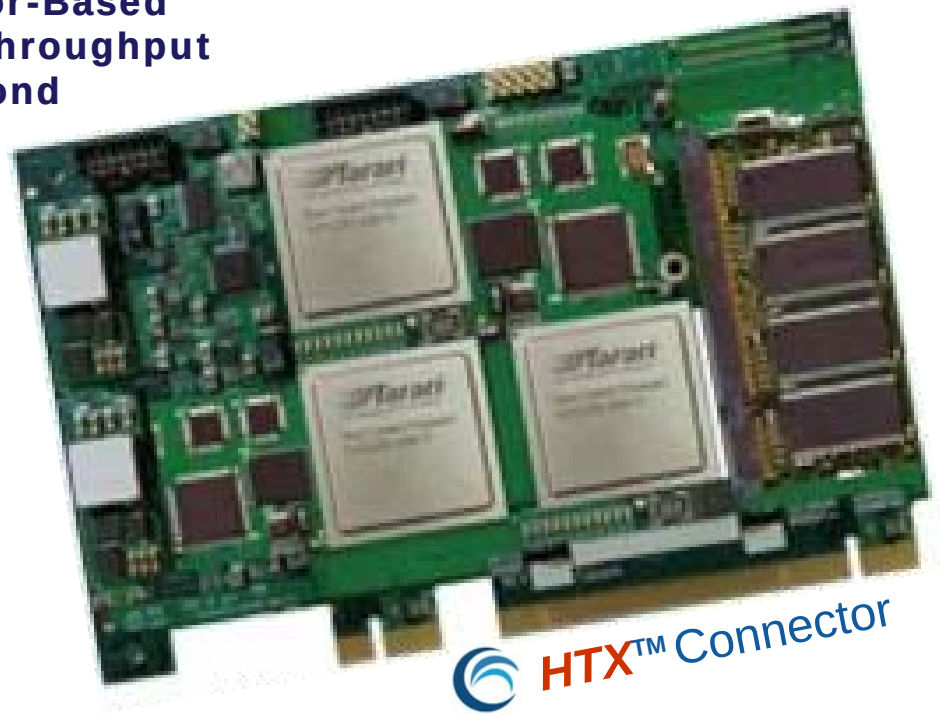

HTX™-Friendly(*) Subsystems (cont.)

TL1550 Virtualization Processor

- From Servers to Expandable Servers
- Virtualizes CPUs, Memory and I/O

(*) HTX Proto/Evaluation Vehicle Available

HTX™-Friendly Subsystems (cont.)

TL1550 Virtualization Processor

- **Cache Coherence Evolution from Bus to MicroNetwork**

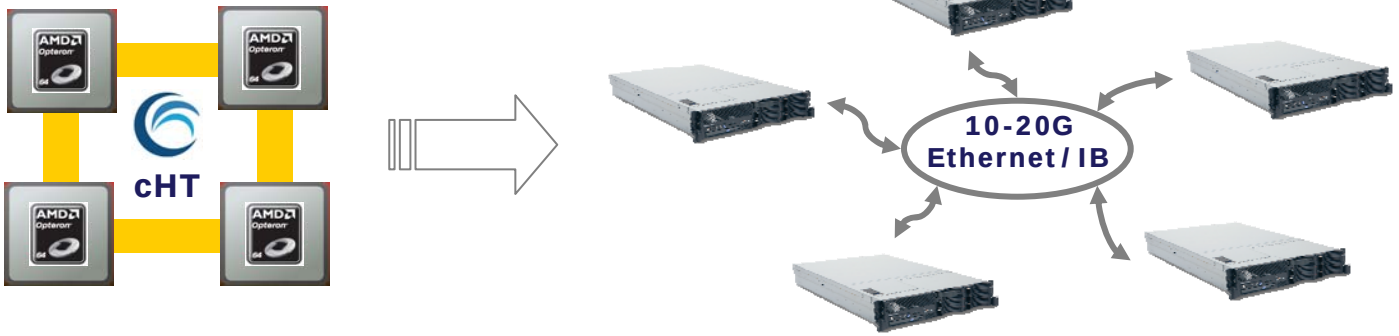

HTX™-Friendly Subsystems (cont.)

TL1550 Virtualization Processor

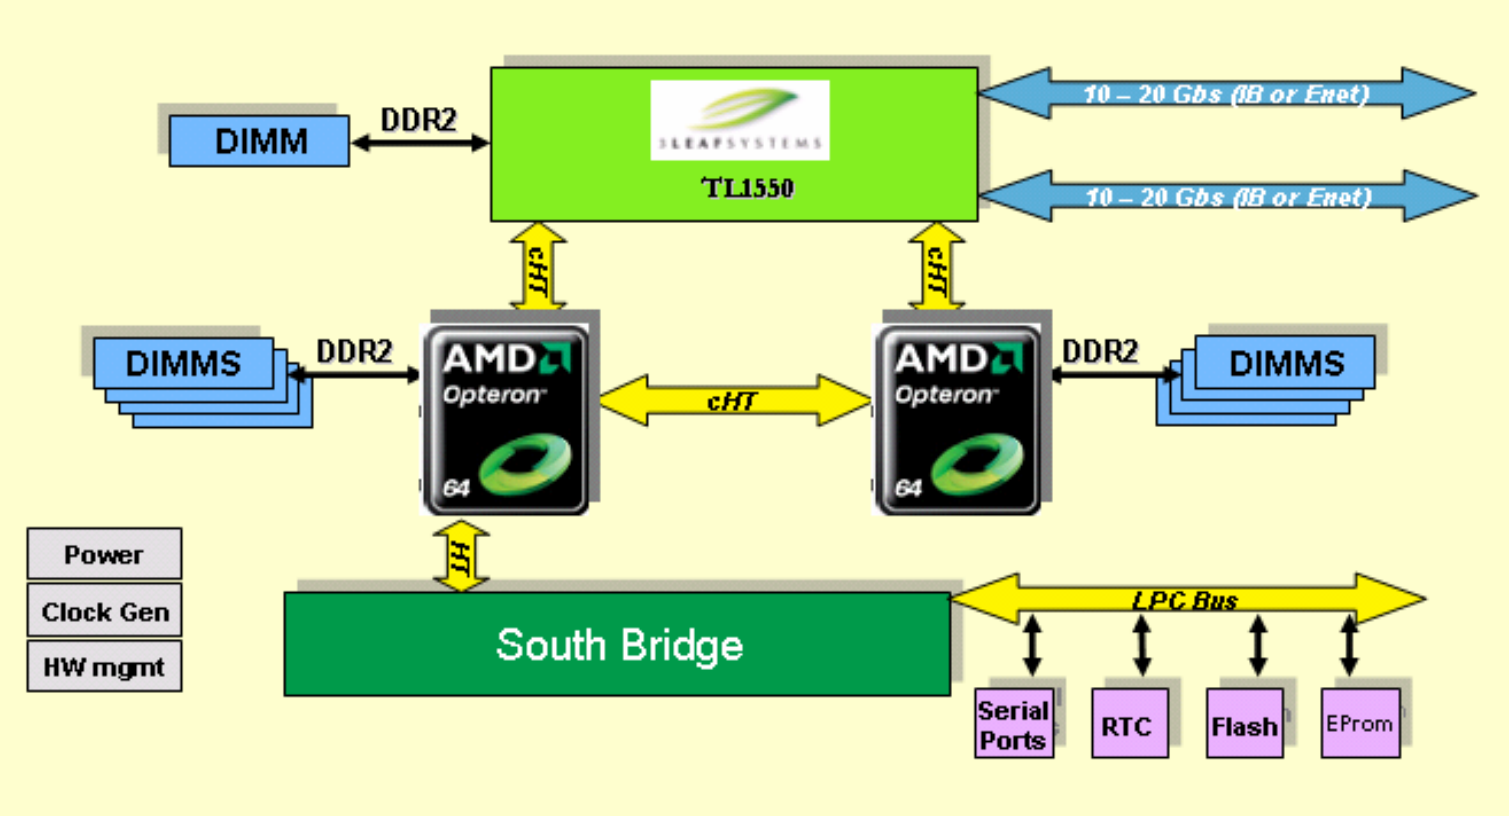

HTX™-Friendly Subsystems (cont.)

3LEAFSYSTEMS

TL1550 Virtualization Processor

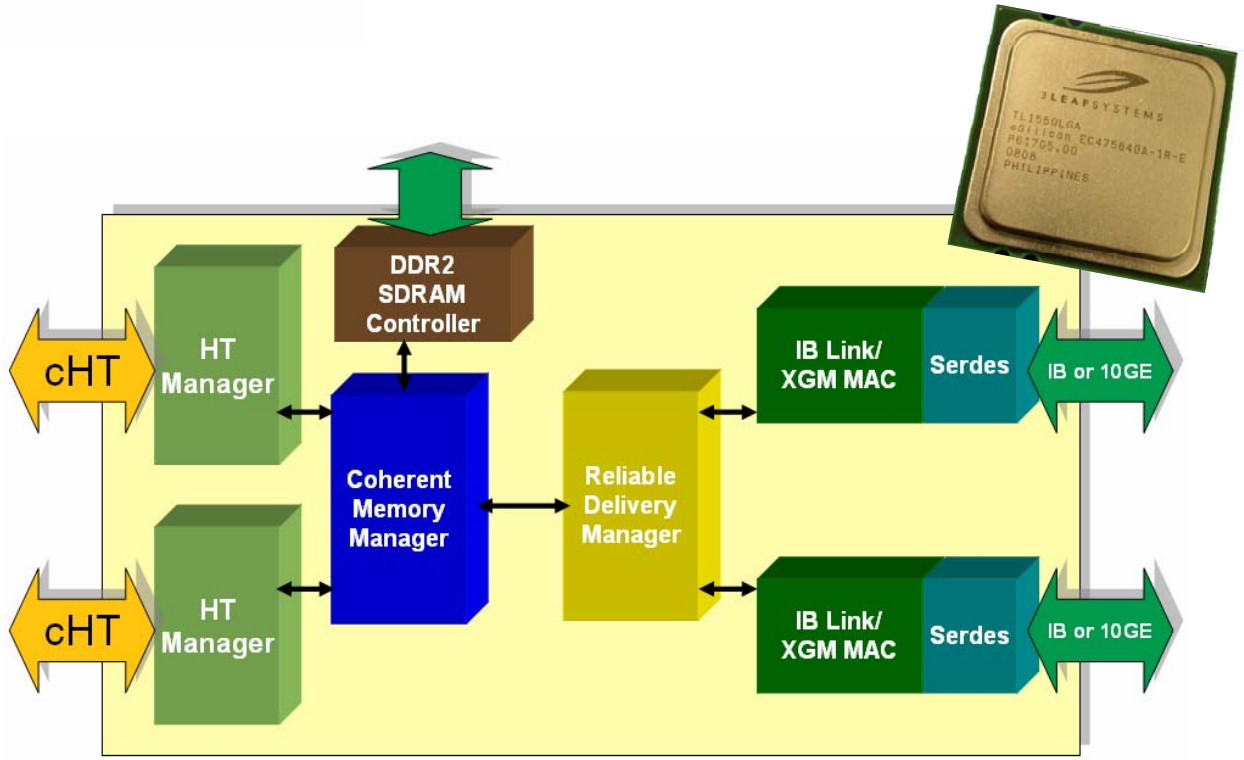

HyperTransport™ CONSORTIUM

Development Support

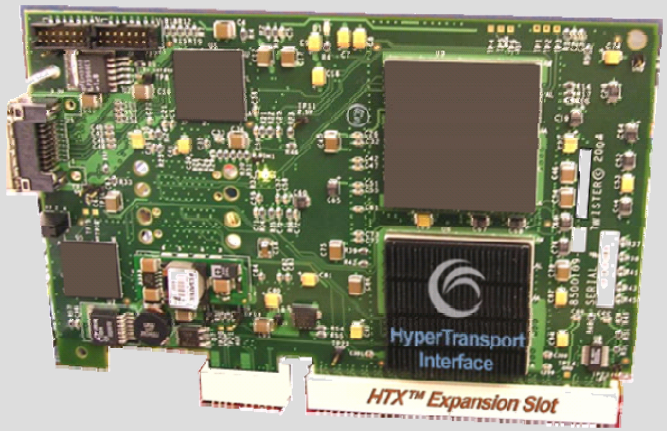

HTX™ Development Support

HyperTransport Center of Excellence Design Support and Services

Universal Board Reference Designs

HTX™

HTX3™

HT and cHT Cores

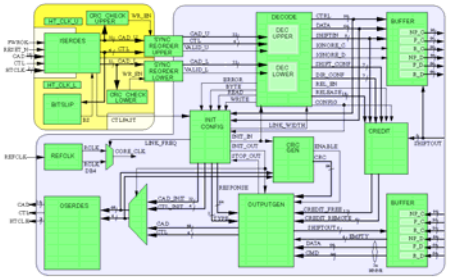

Test Tools

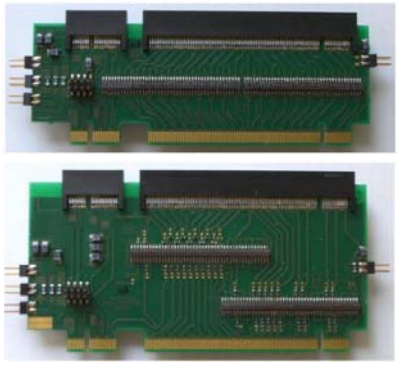

Design Support Product Validation

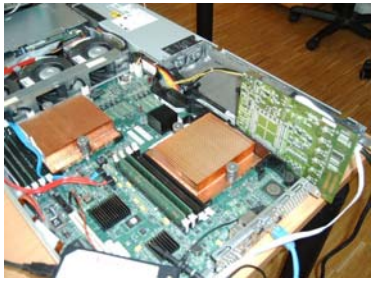

HTX™ Development Support (cont.)

HTX ™ 10GE NIC Design Kit

- Validated and Tested 10 Gigabit Ethernet Solution
- Dowloadable Free-of-Charge from HTC Web Site

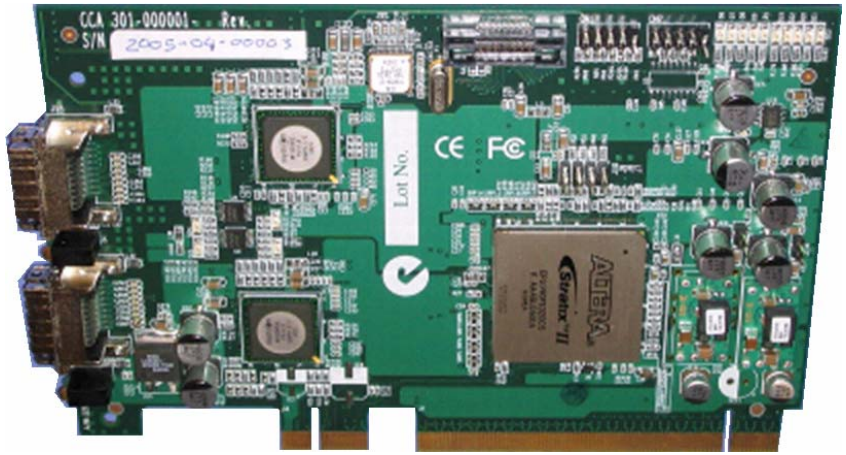

HTX™ Connector

- 2x 10GE Ports
- Stratix II FPGA

Fast Time-to-Market

HTX™ Development Support (cont.)

HTX ™ FPGA Development Board

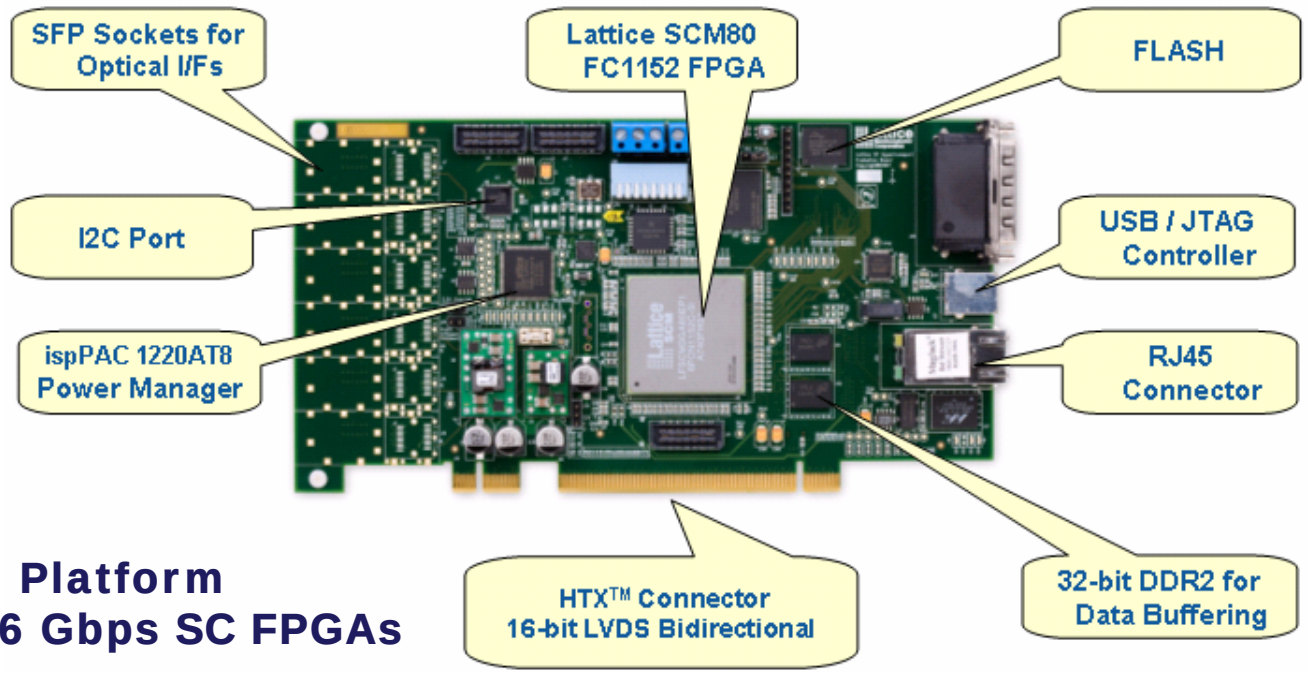

Rapid Design Platform for Lattice 1.6 Gbps SC FPGAs

HTX™ Connector

HyperTransport™ CONSORTIUM

BIOS Support

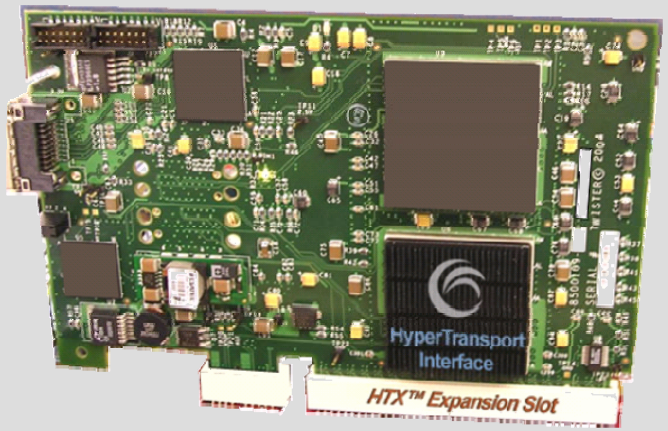

HTX™ BIOS Support

Consortium-Defined and Standardized *HTX*™ BIOS Specifications

Cross-Compatibility

Fast Time-to-Market

HyperTransport™ CONSORTIUM

Why HTX?

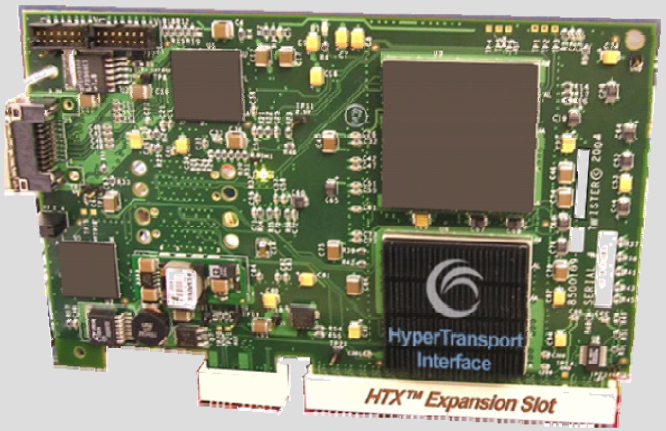

Unique Technical Values

HTX™ Provides Superior Performance and Architectural Capabilities In Line with HPC Market Expectations

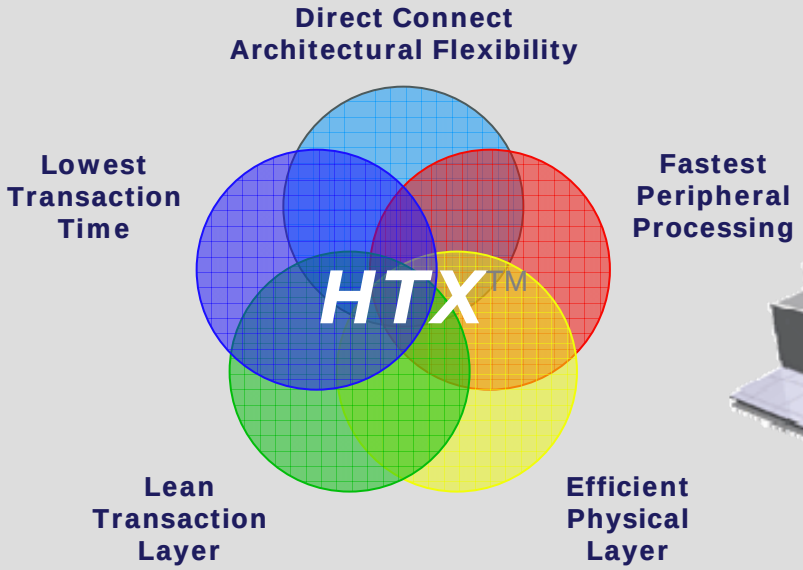

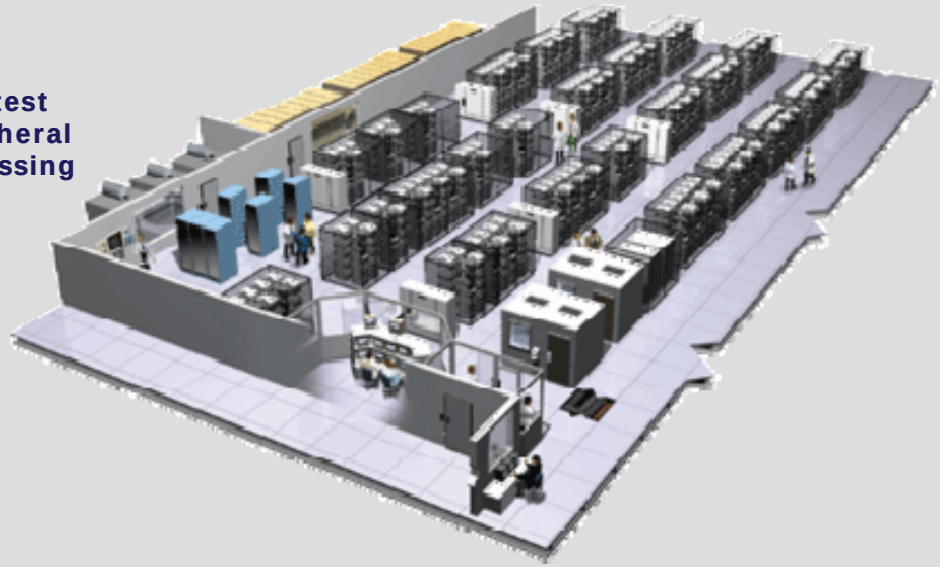

HyperTransport™ CONSORTIUM

HTX™ Complements PCI Express

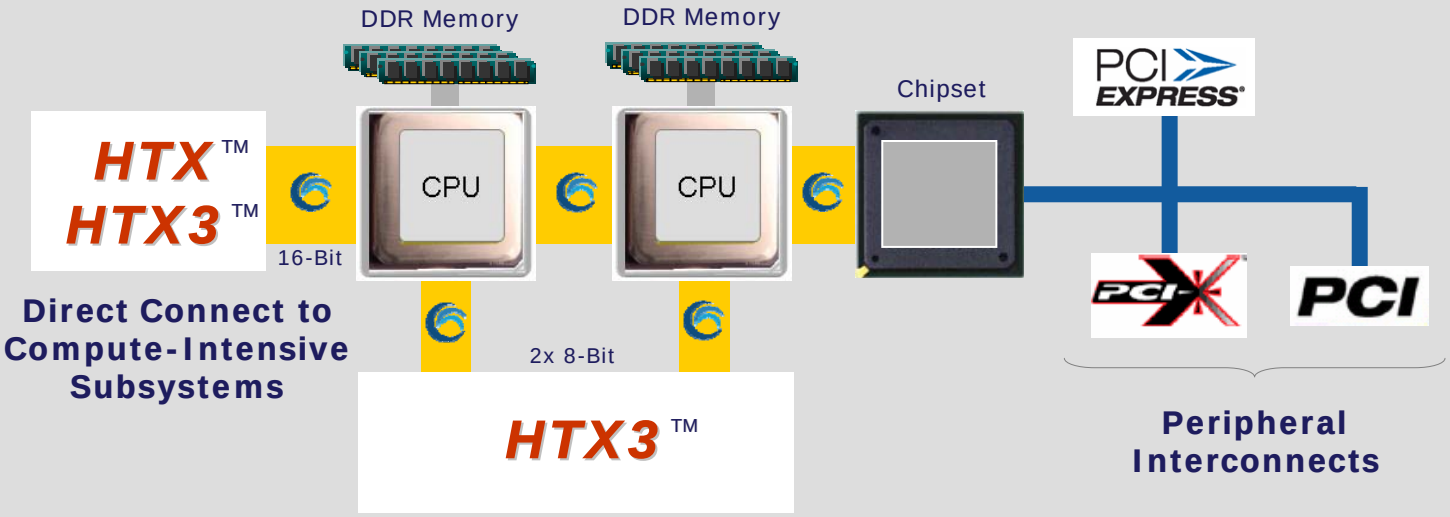

With its Unique Architectural and Performance Edge over General Purpose Interconnects, HTX is an Ideal Complement to PCI-Class Slot Connector Interfaces

HTX™ Business Values

HTX™ is a Competitive Differentiator for High-Performance Total Solution Vendors

- **Negligible BOM Cost**
  - Leverages PCIe Economy of Scale
- **System Flexibility**
  - Single MB Serves Multiple Markets
- **Growing HTX Subsystem Portfolio**
  - Coprocessors, Accelerated 10GE, ccNUMA Cluster Controllers, Content Processors, Others Released Soon
- **Up-Selling Enabler**
  - Delivers Greater End-User Value

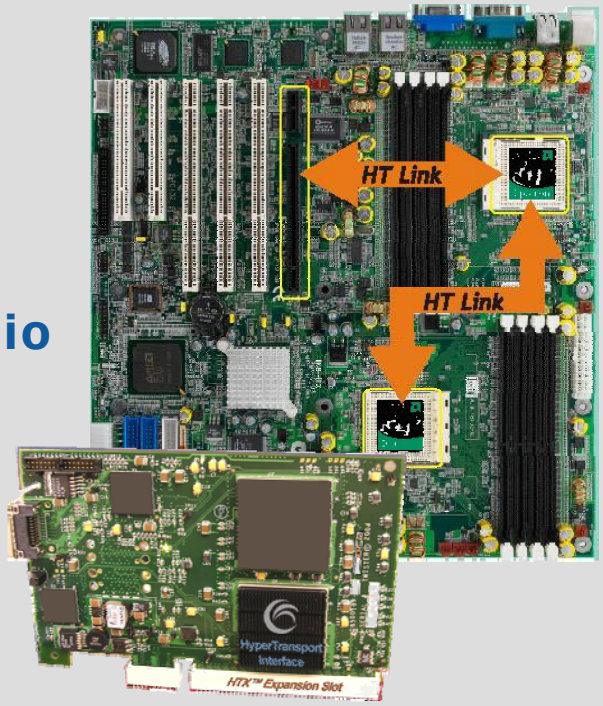

HyperTransport™ CONSORTIUM

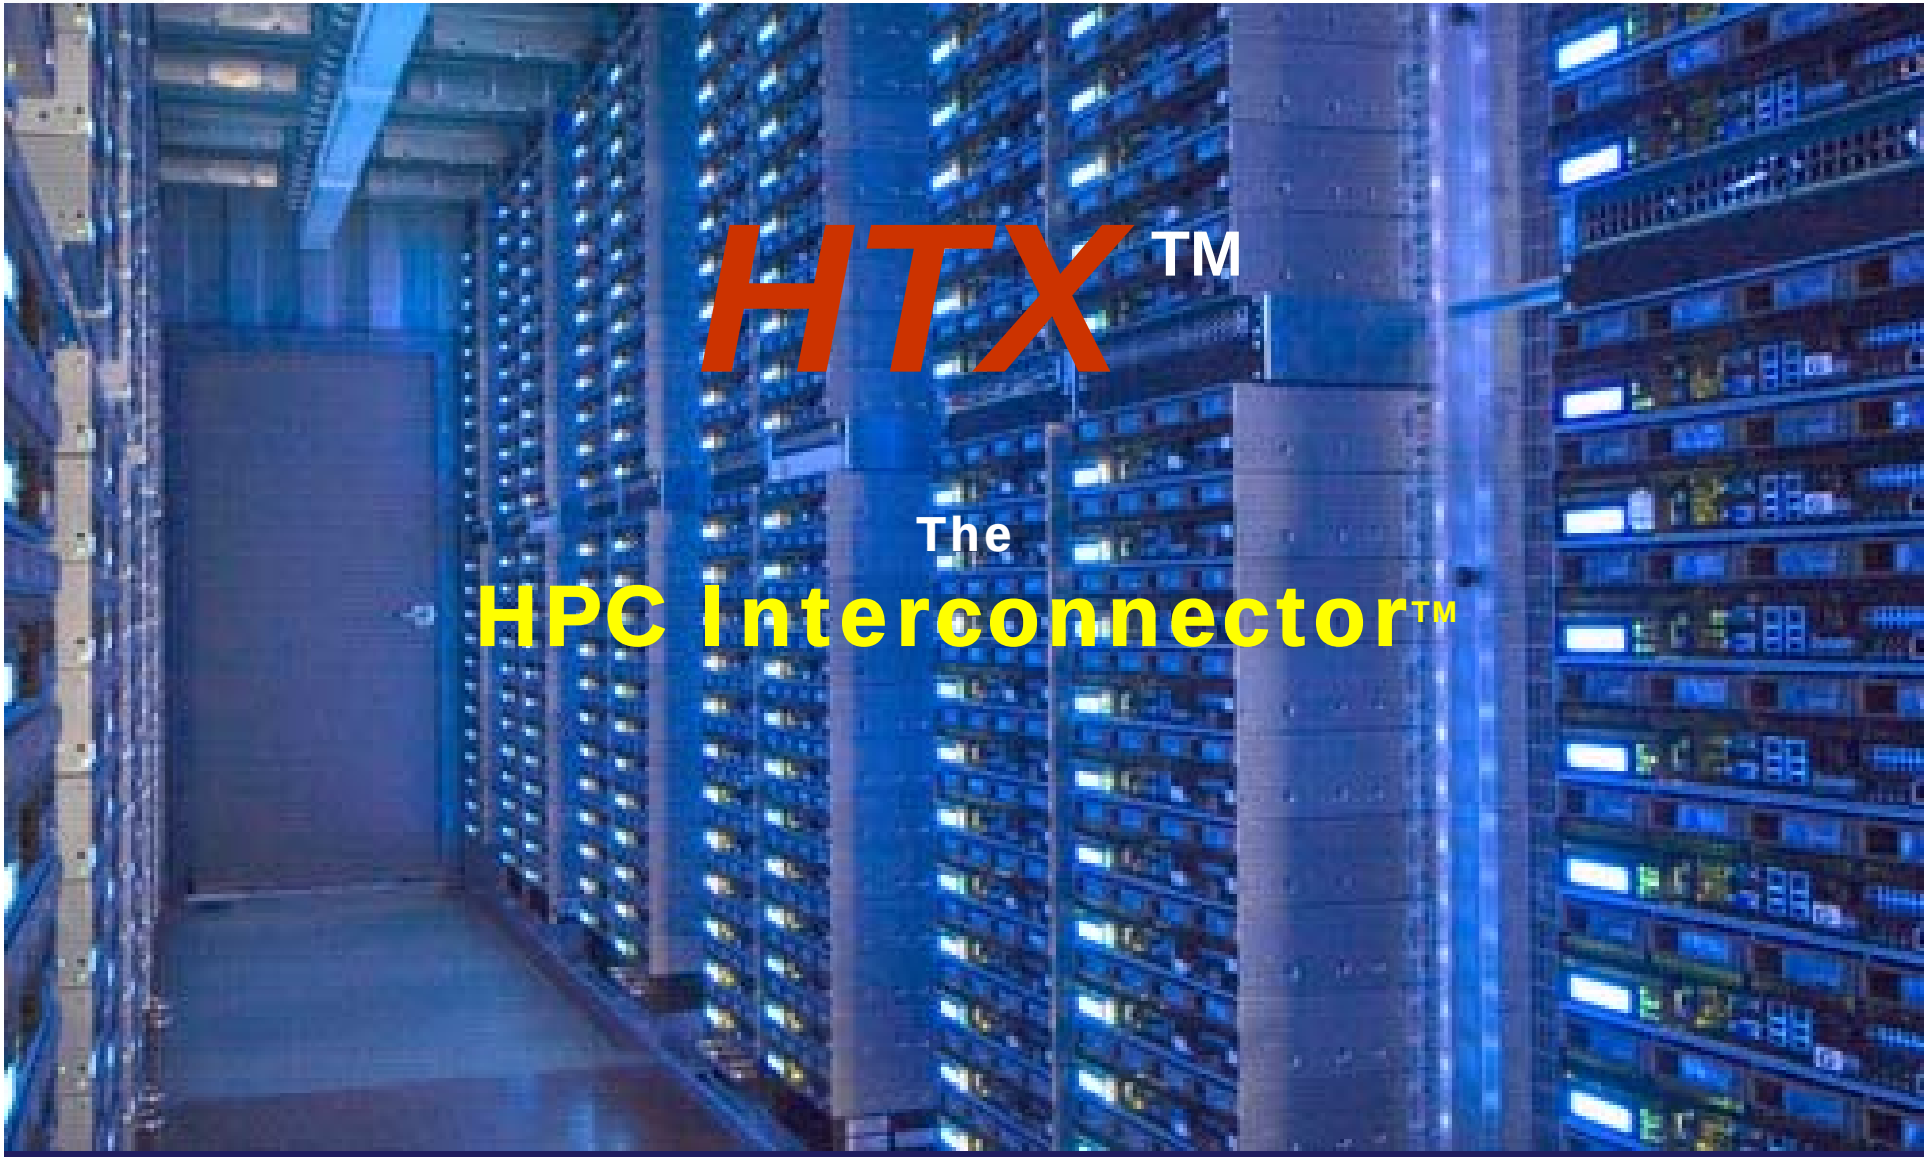

HTX™

The HPC Interconnector™

Can Interconnect Technologies Meet the HPC Challenge?

HyperTransport™ Can With Strong Core Values

- Architectural Simplicity
- Protocol Efficiency
- Full Scalability
- Low Design Cost
- Extensive Support
- Vast Product Ecosystem

HyperTransport™ CONSORTIUM

With HyperTransport™ The Technology Industry Wins