# Application of General Purpose HPC Systems in HPEC

David Alexander
Director of Engineering
SGI, Inc.

Igniting Innovation and Leadership

# SGI HPC System Architecture

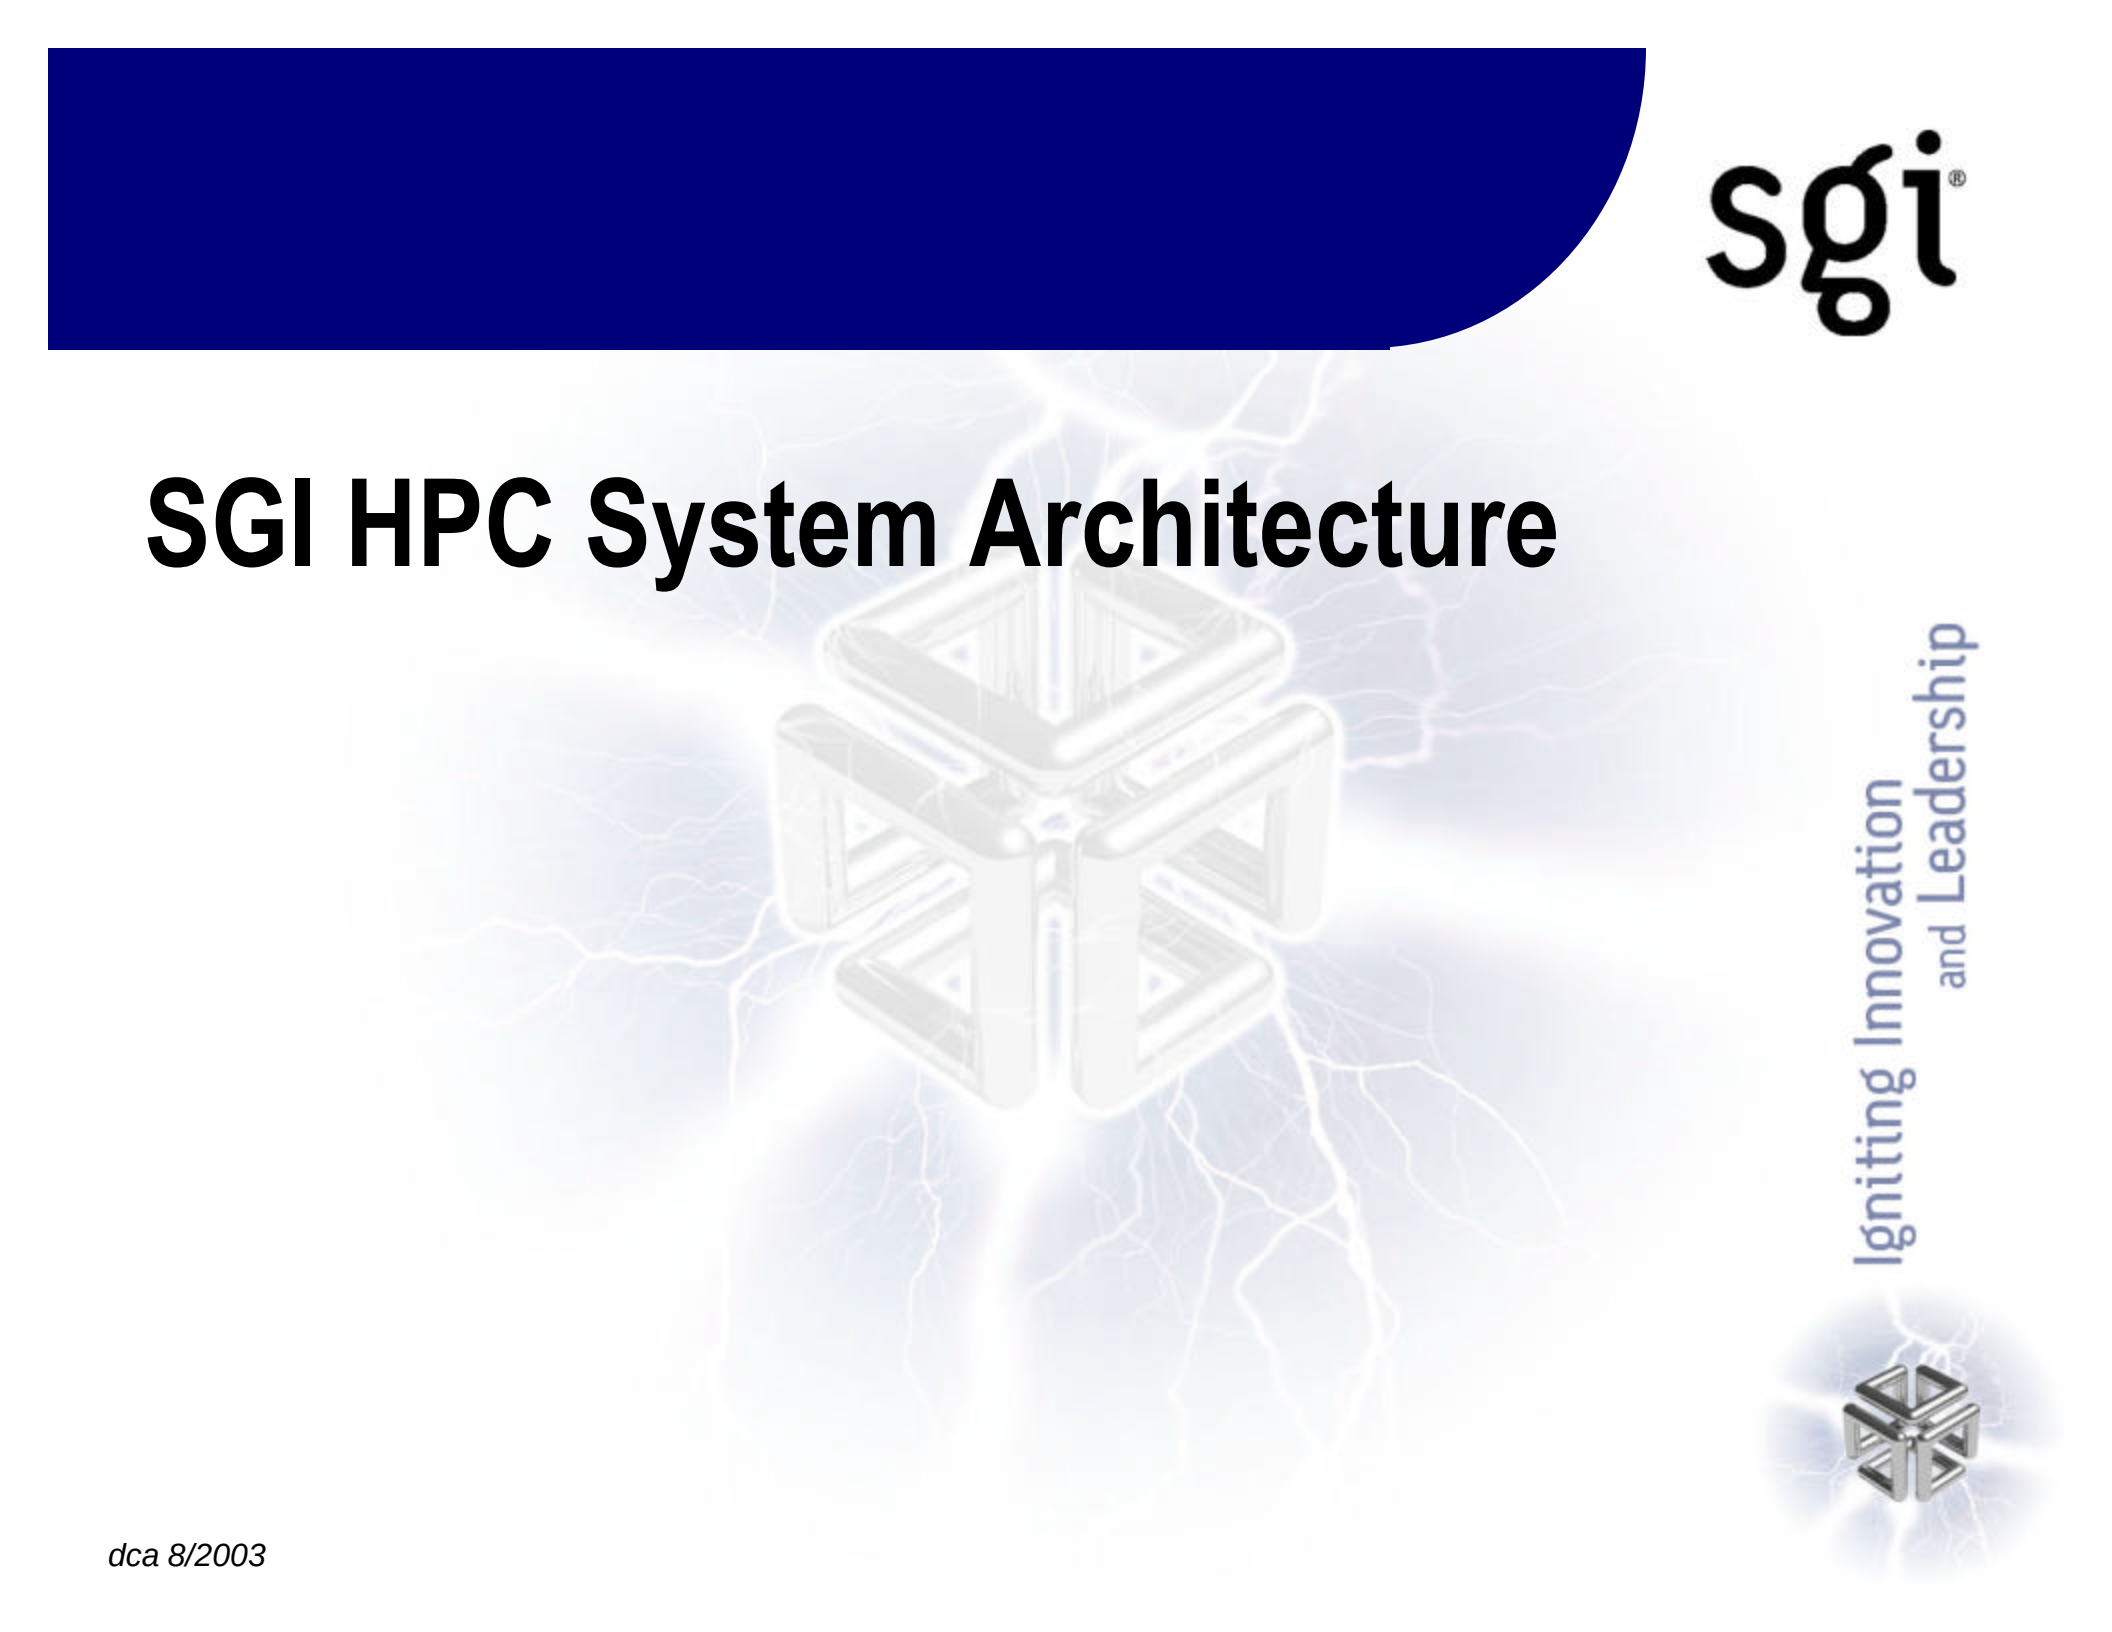

Igniting Innovation and Leadership

*dca 8/2003*

# SGI Scalable ccNUMA Architecture
## Basic Node Structure and Interconnect

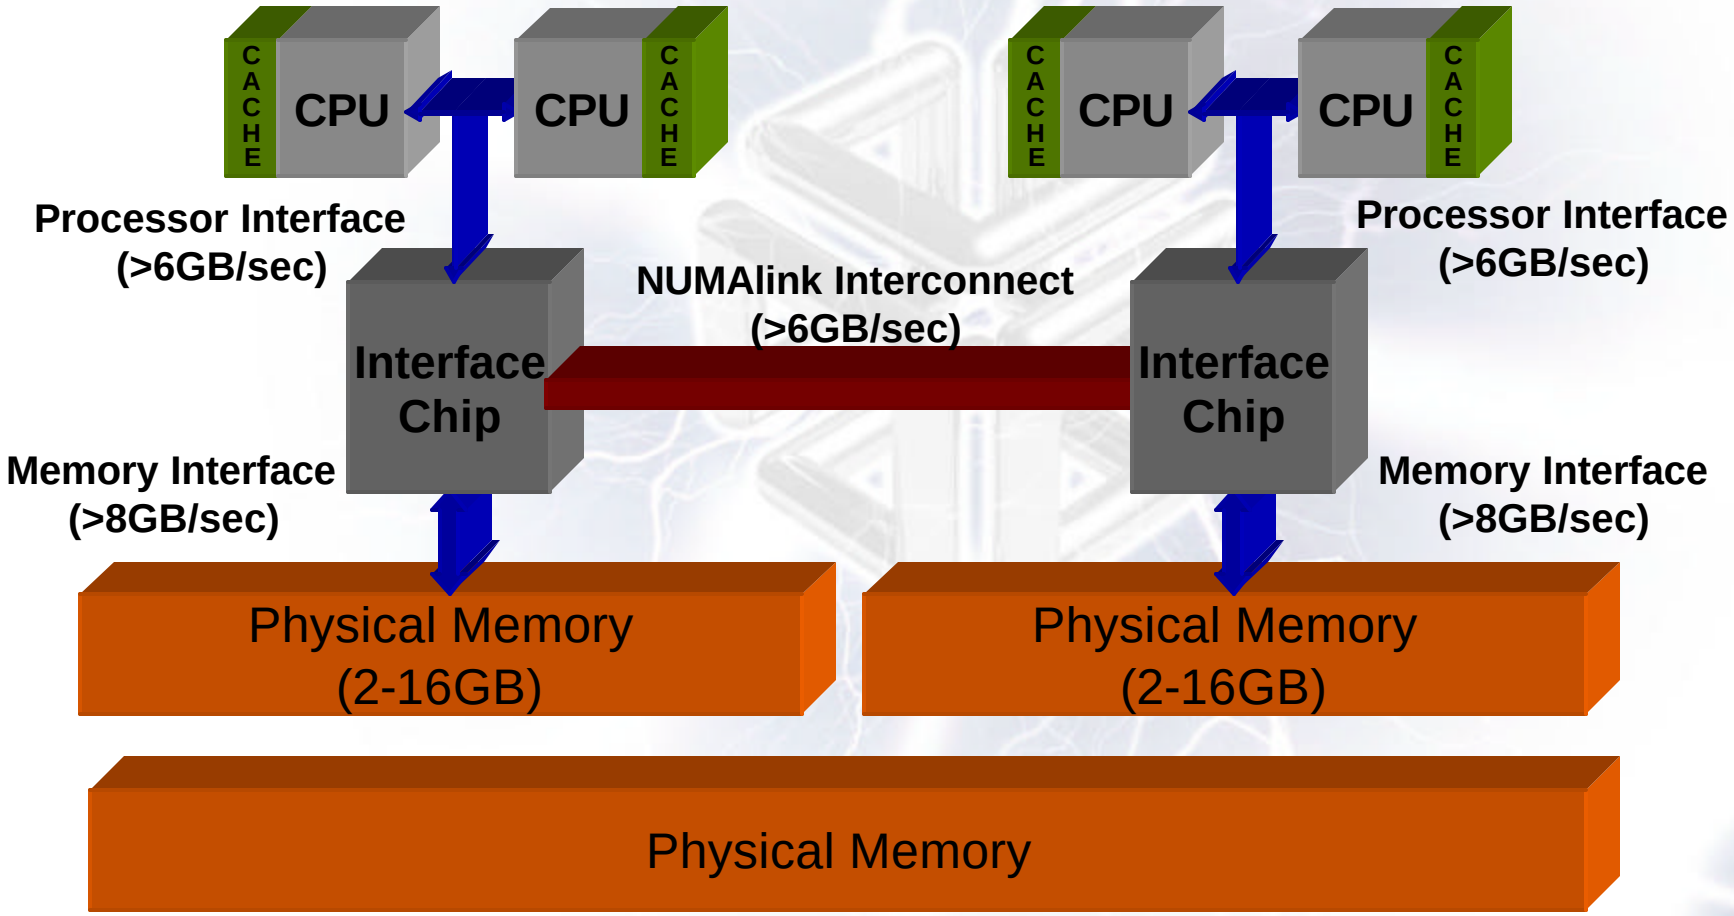


Igniting Innovation and Leadership

# SGI Scalable ccNUMA Architecture
## Scaling to Large Node Counts

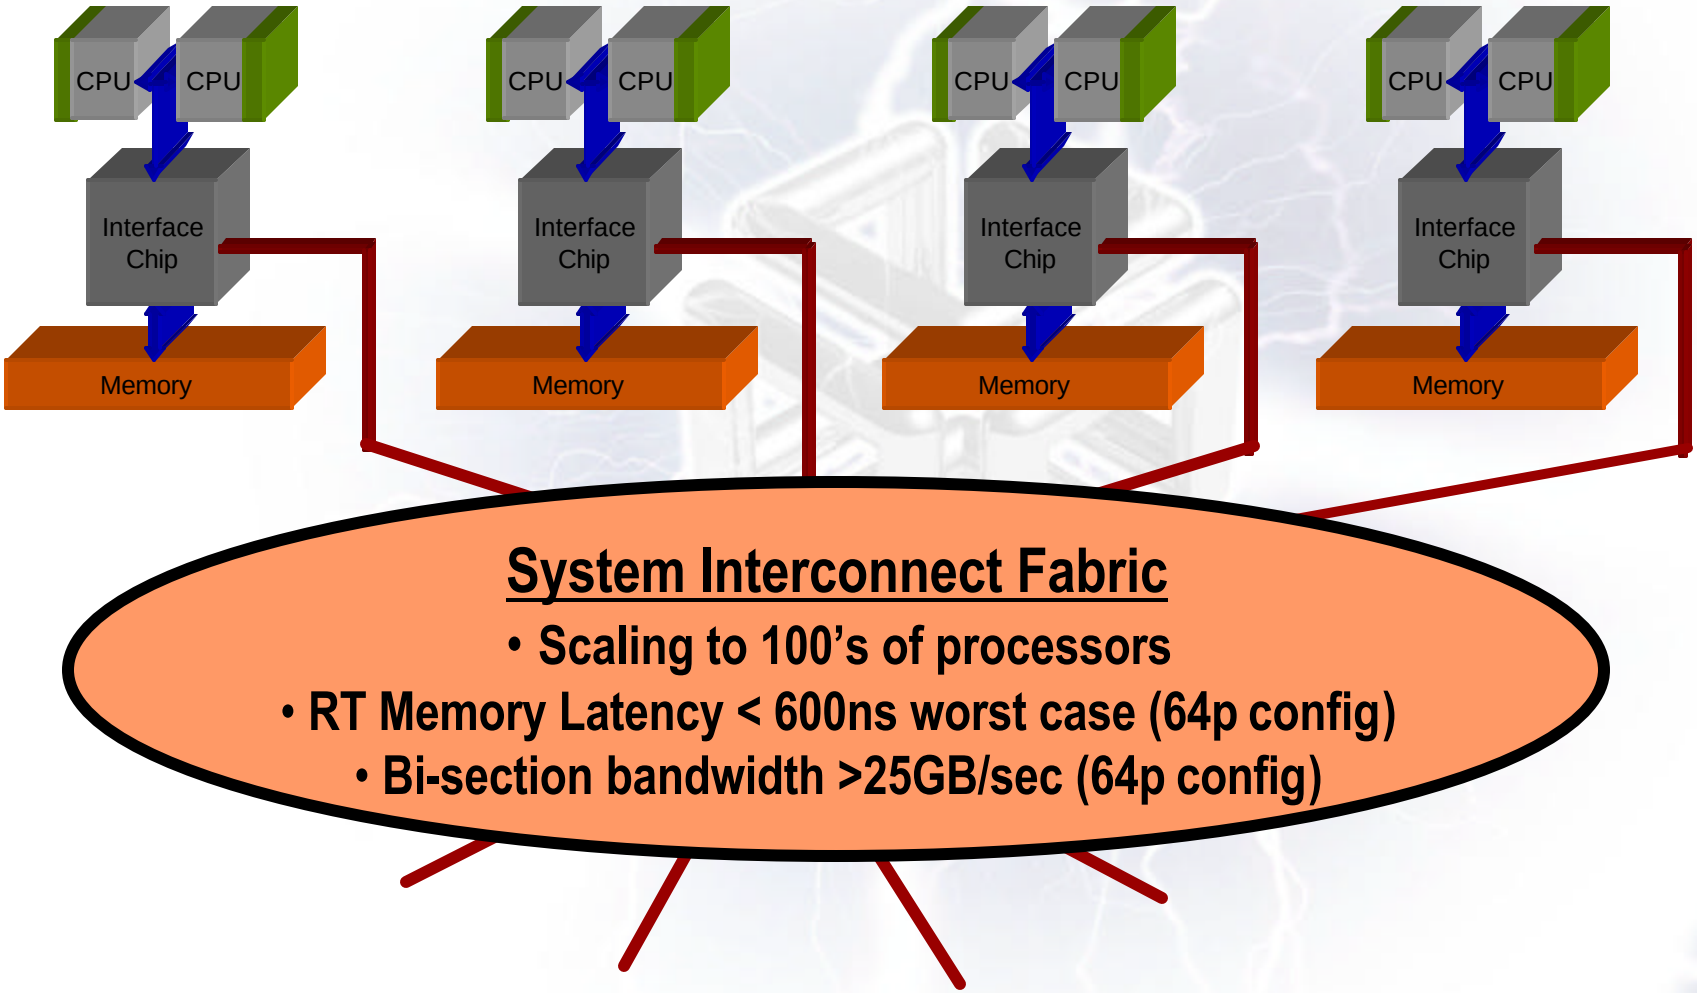


Igniting Innovation and Leadership

# SGI Scalable ccNUMA Architecture

## SGI® NUMAflex™ Modular Design

System Interconnect

CPU & Memory

System I/O

Standard I/O Expansion

High BW I/O Expansion

Graphics Expansion

Storage Expansion

Modules from SGI® Onyx® 3000 Series

**Performance:** High-bandwidth interconnect with very low latency

**Flexibility:** Tailored configurations for different dimensions of scalability

**Investment protection:** Add new technologies as they evolve

**Scalability:** No central bus or switch; just modules and NUMAlink™ cables

Igniting Innovation and Leadership

# SGI Scalable ccNUMA Architecture
## The Benefits of Shared Memory

sgi®

### Traditional Clusters

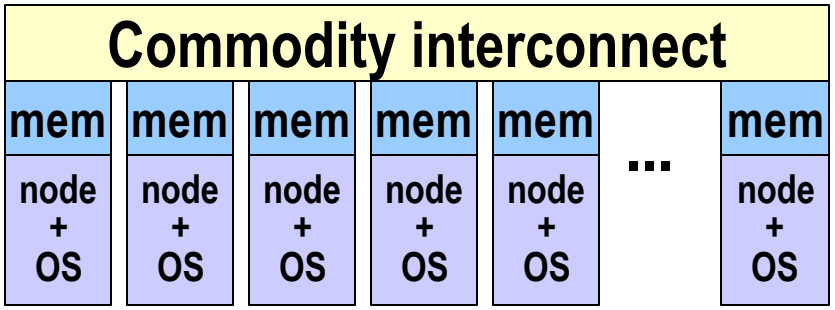


### SGI® Altix™ 3000

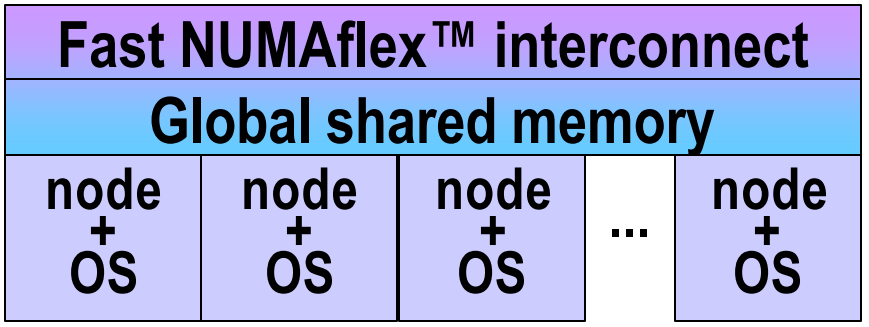


## What is shared memory?

- All nodes operate on one large shared memory space, instead of each node having its own small memory space

## Shared memory is high-performance

- All nodes can access one large memory space efficiently, so complex communication and data passing between nodes aren't needed
- Big data sets fit entirely in memory; less disk I/O is needed

## Shared memory is cost-effective and easy to deploy

- The SGI Altix 3000 family supports all major parallel programming models
- It requires less memory per node, because large problems can be solved in big shared memory
- Simpler programming means lower tuning and maintenance costs

# SGI® Altix™ 3000 HPC Product

Igniting Innovation and Leadership

*dca 8/2003*

# SGI® Altix™ 3000
## Overview

- First Linux® node with 64 CPUs in single-OS image
- First clusters with global shared memory across multiple nodes
- First Linux solution with HPC system- and data-management tools
- World-record performance for floating-point calculations, memory performance, I/O bandwidth, and real technical applications

Igniting Innovation and Leadership

# SGI® Altix™ 3000
## Fusion of Powerhouse Technologies

sgi®

**SGI® supercomputing technology, Intel's most advanced processors, and open-source computing**

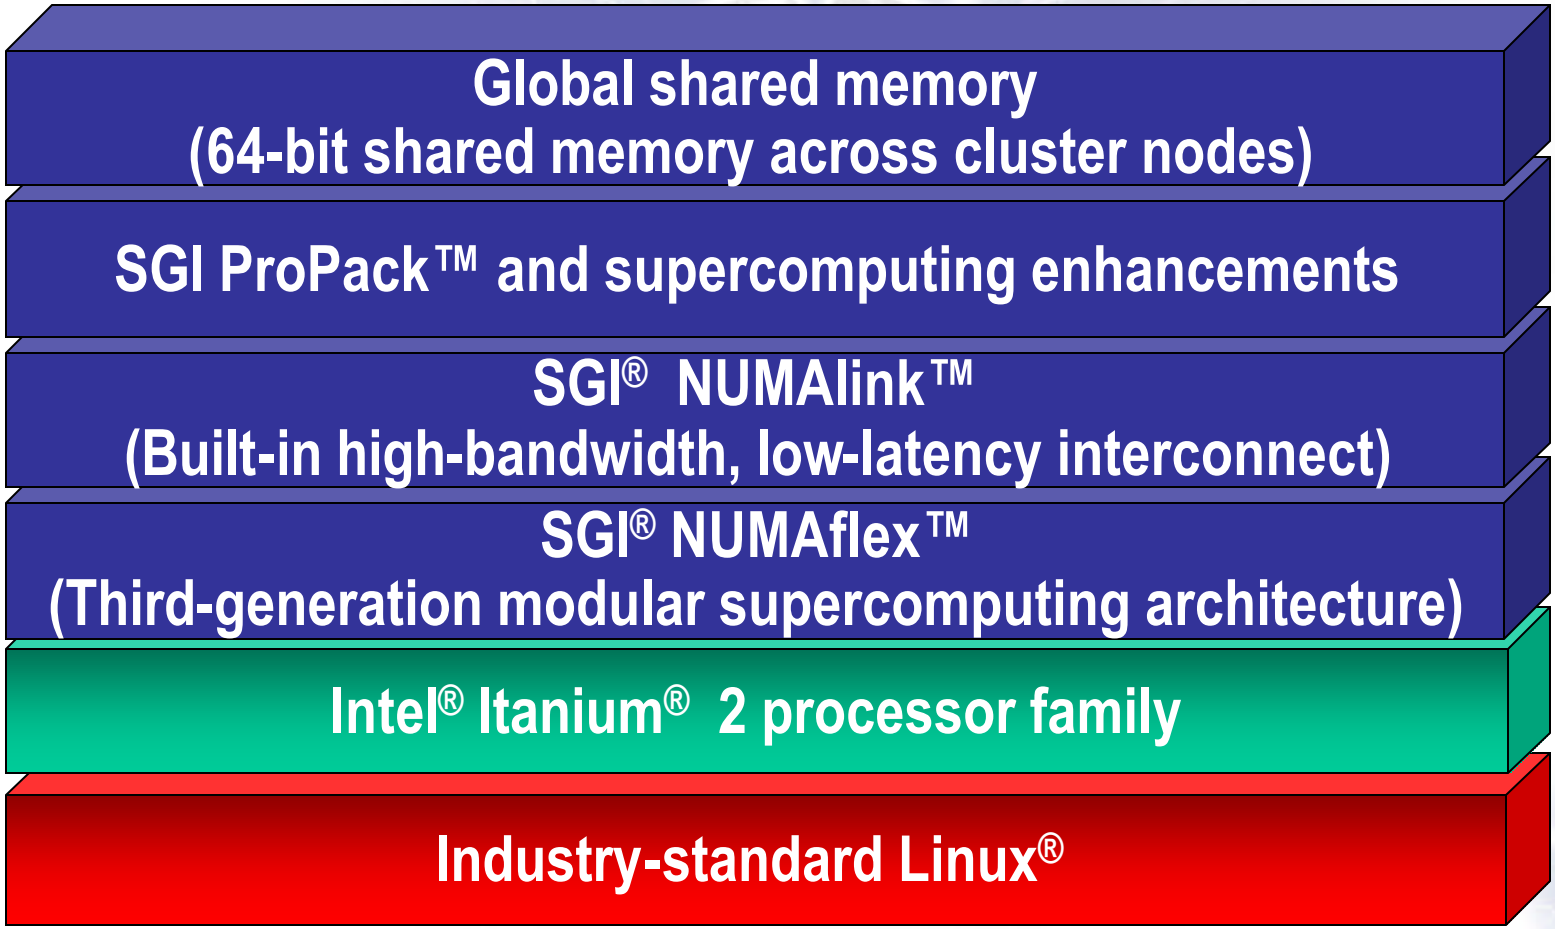


Igniting Innovation and Leadership

# SGI® Altix™ 3000 Family
## Extreme Power, Extreme Potential

**Model 3300 Servers**

Single-node entry offering
4–12 1.30 GHz Itanium® 2 processors, 3MB L3 cache
Up to 96GB memory

**Model 3700 Superclusters**

Scalable performance offering
4–64 Itanium 2 processors in a single node
- 1.30 GHz/3MB L3 cache
- 1.50 GHz/6MB L3 cache

Shared memory across nodes
Scalable to 2,048 processors, 16TB memory
Nodes up to 64P, 4TB memory

Igniting Innovation and Leadership

# Multi-Paradigm Architecture

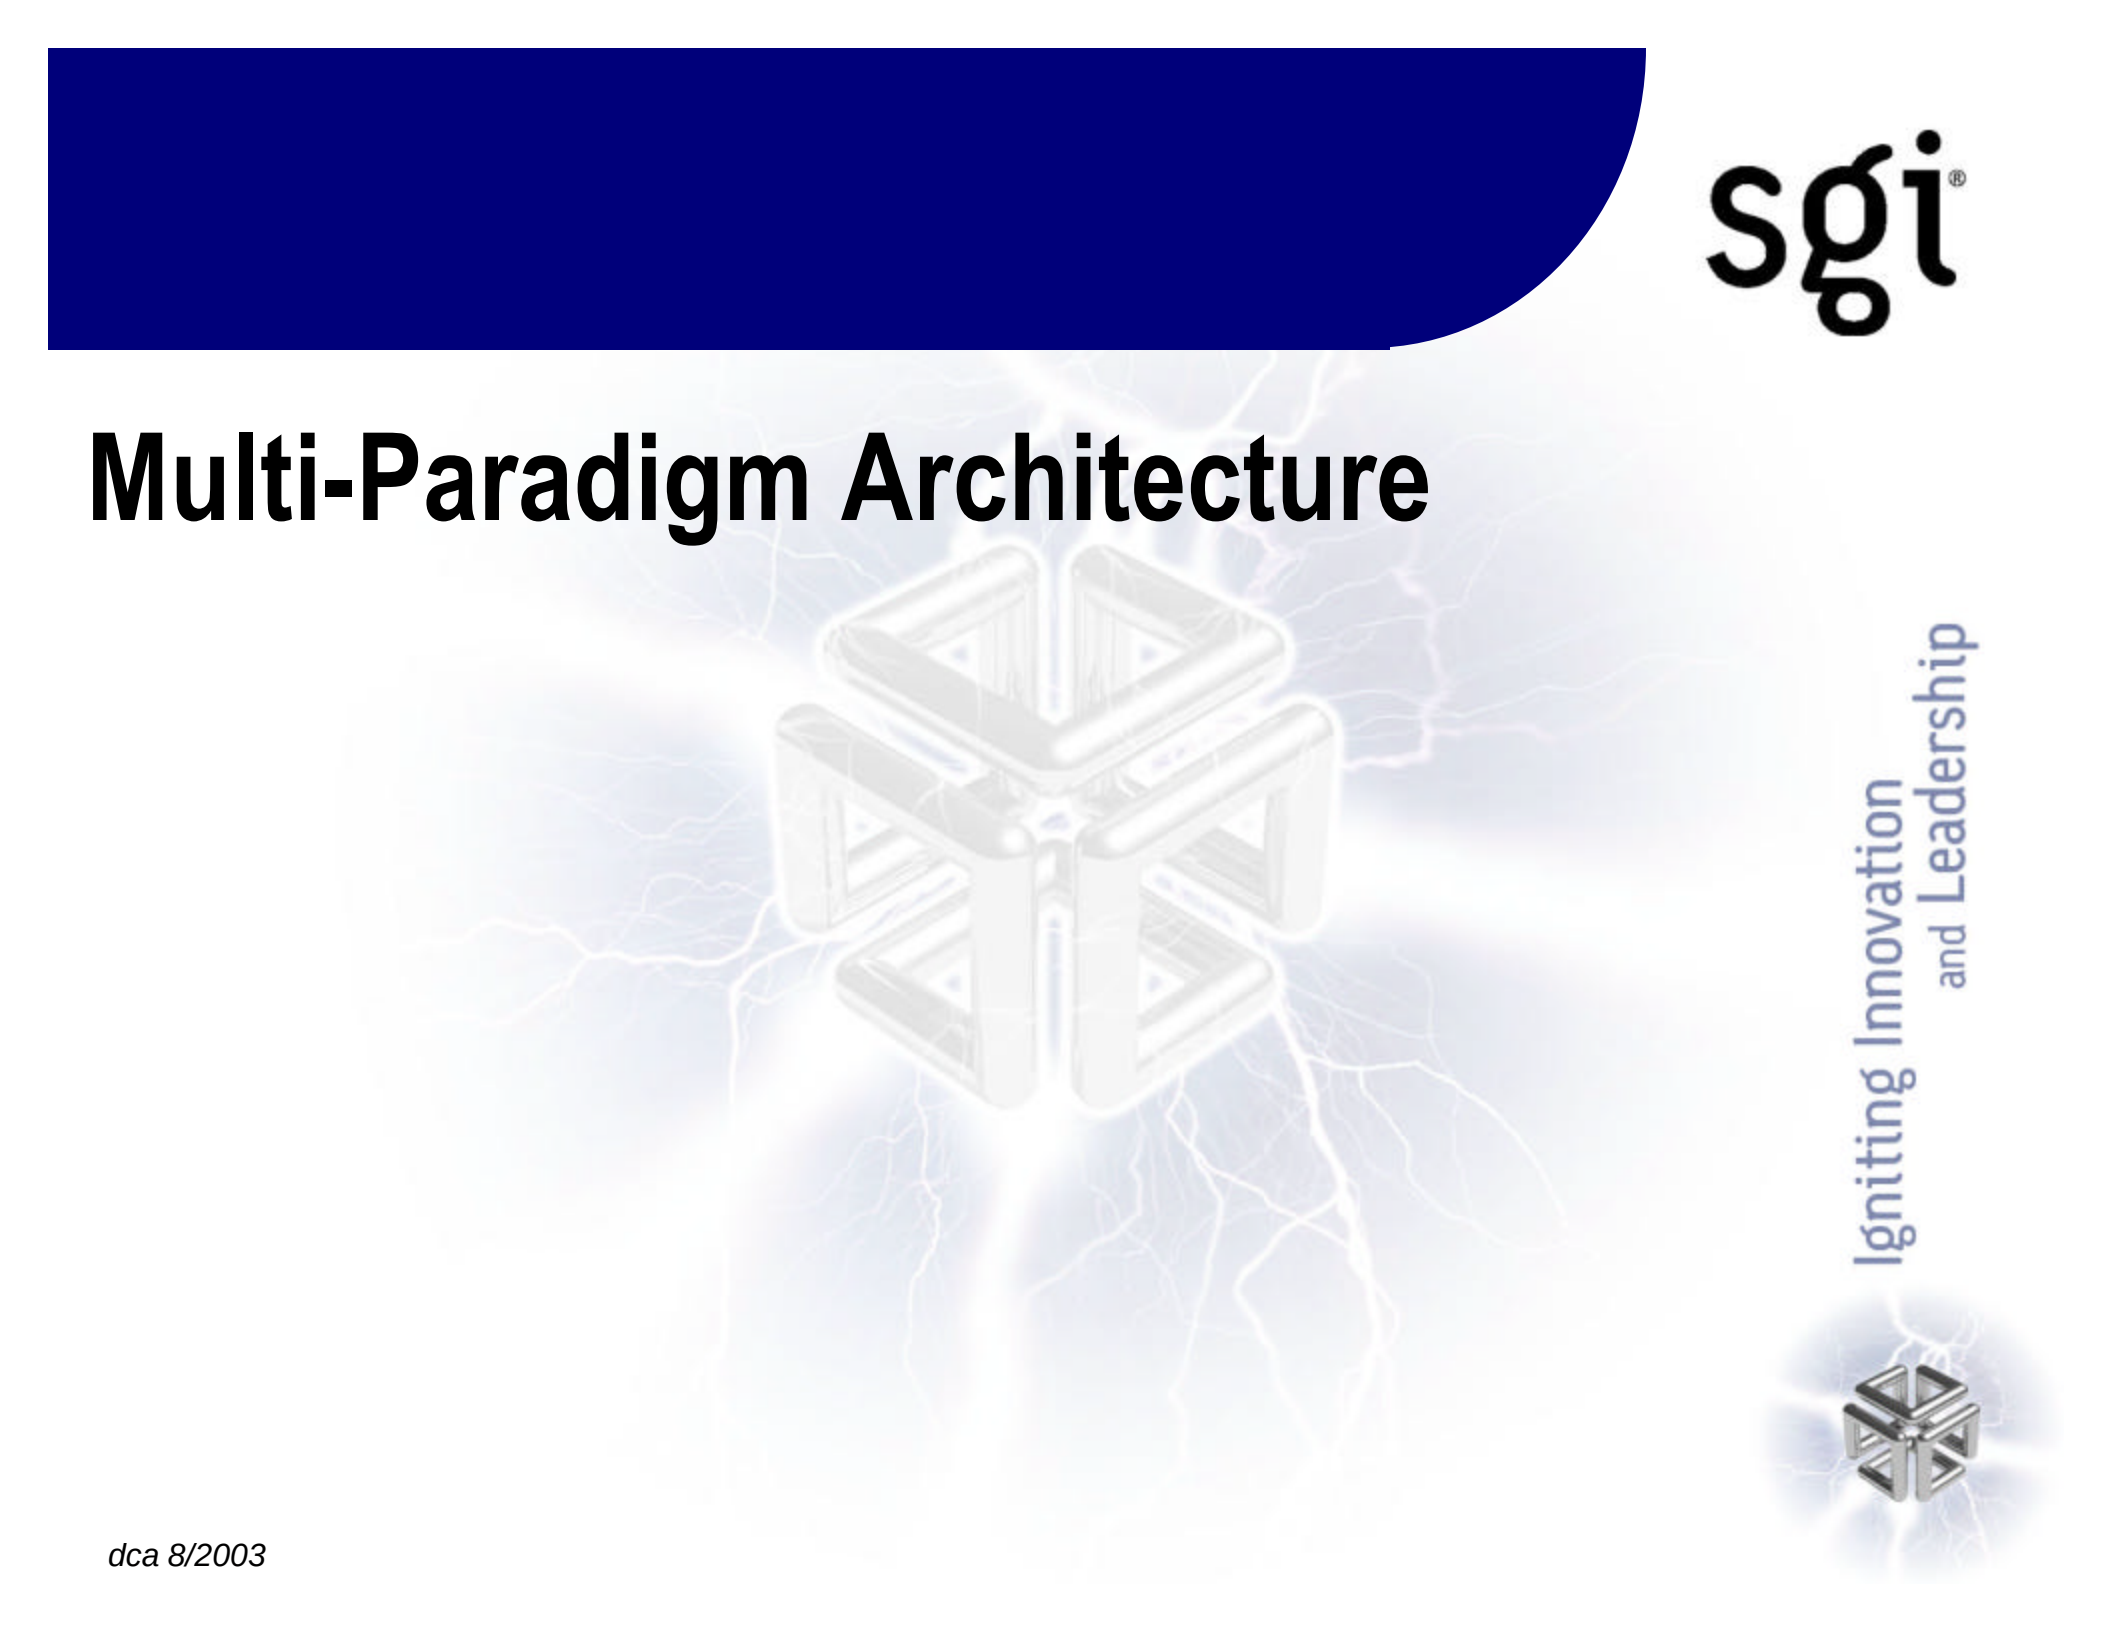

Igniting Innovation and Leadership

# Multi-Paradigm Architecture
## Overview

- R • NUMAlink system interconnect
- C • General-purpose compute nodes
- IO • Peer-attached general purpose I/O
- MSE • Mission-specific accel. and/or I/O (MSE)
- V • Integrated graphics/visualization

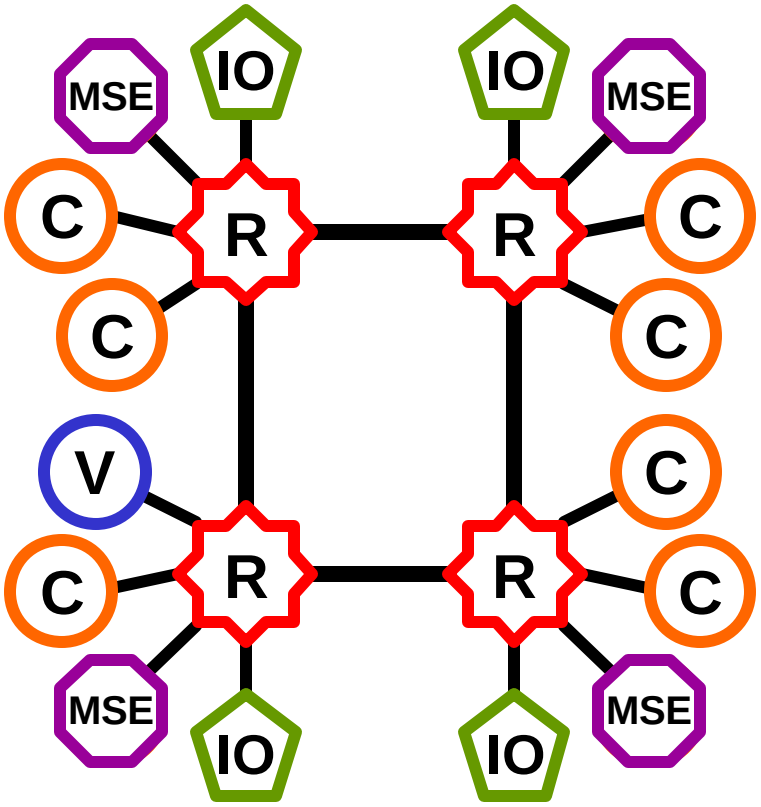

# Multi-Paradigm Architecture
## Mission Specific Element (MSE)

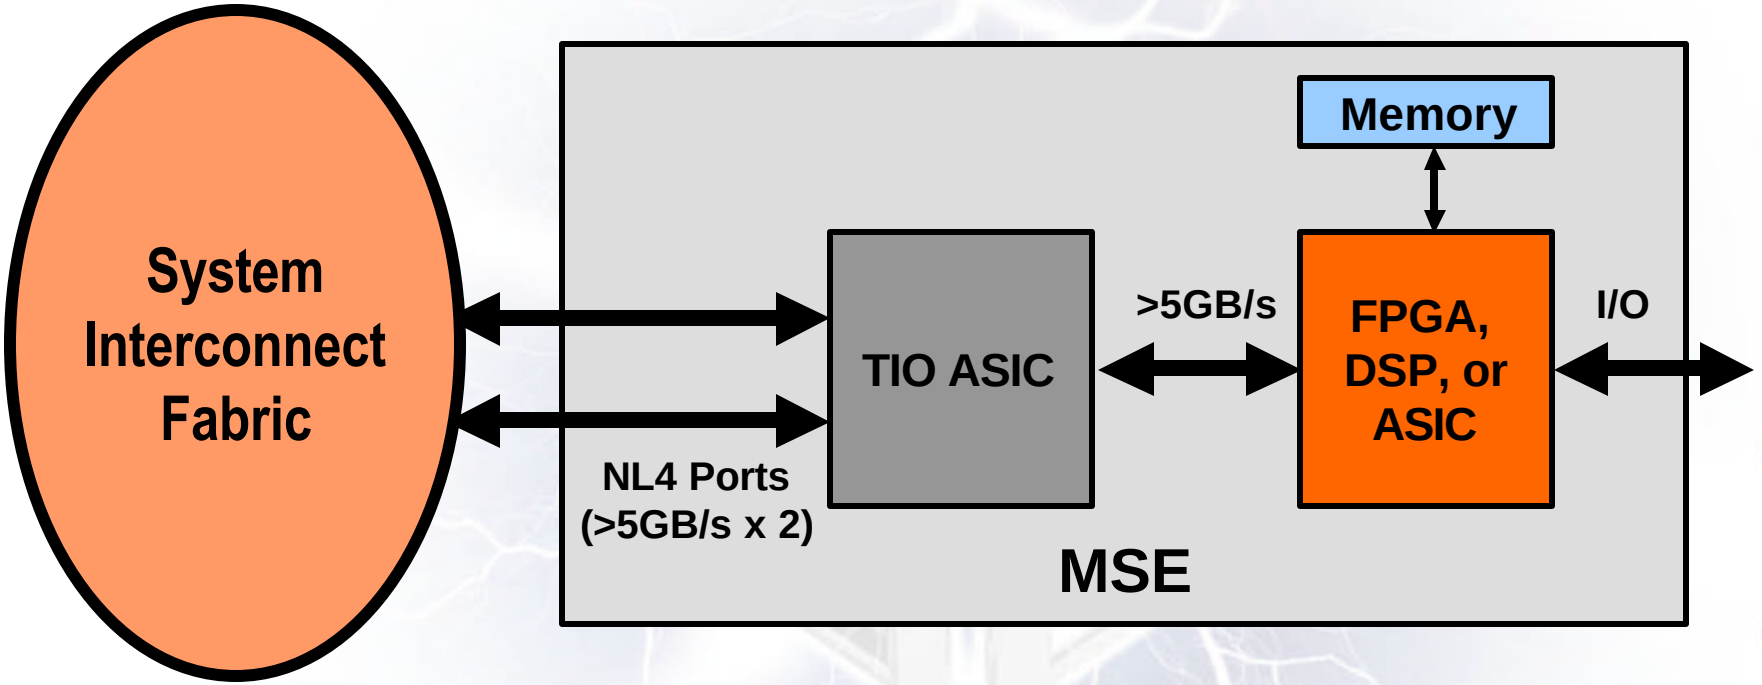


Igniting Innovation and Leadership

## Implementation variants:

- Customer supplied/loaded FPGA algorithms and/or ASIC/DSP
- Subroutine or standard library acceleration
- Specific use “appliance”

# Embedded High Performance Computing (EHPC)

Igniting Innovation and Leadership

*dca 8/2003*

# Embedded High Performance Computing
## Unified Development/Deployment Environment

## EHPC Today

Different Development and Deployment Environments

**Development Environment**
**HPC methods, tools, hardware**
**Architecture provides freedom to develop new approaches**

Platform Port
Months of work
Millions of dollars
Significant schedule risk
Significant architectural/performance risk

**Deployment**
**Proprietary hardware and software**
**"Islands of memory" narrows algorithm options**

## EHPC Vision

Unified Development and Deployment

**Enhanced software development productivity**
**Superior performance and HW utilization**
**Demonstration to deployment in days**
**Benefits from mainstream HPC advances**

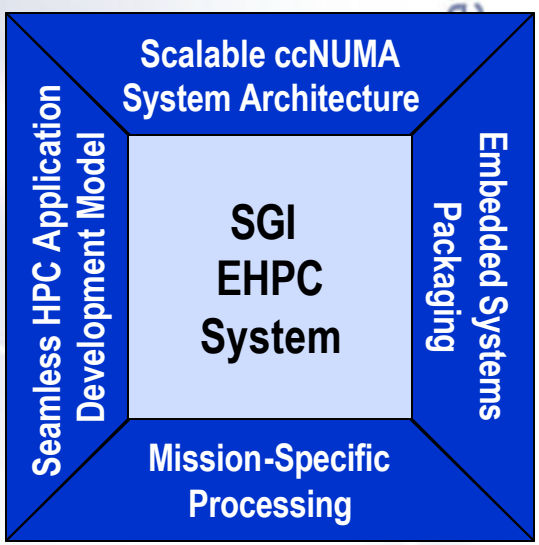

# Embedded High Performance Computing
## EHPC Platform

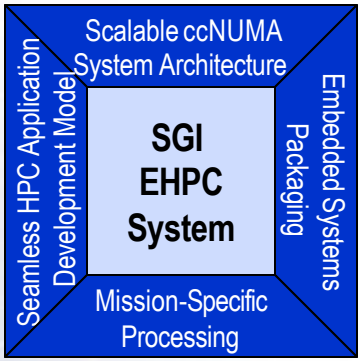


- Unified software environment for development, prototype and deployment
- Full multi-paradigm computing support
- Field-upgradeable mission specific processing acceleration
- Highly optimized performance per watt and performance per slot
- Design to fit into established environmentals, form factors, and interfaces
  - 6U Eurocard form factor
  - Passive, slot configurable backplane
  - PMC module connectivity to standard interfaces
  - Able to address oceanic, ground and airborne environmentals

Igniting Innovation and Leadership

# Embedded High Performance Computing
## Family of 6U Form Factor Modules

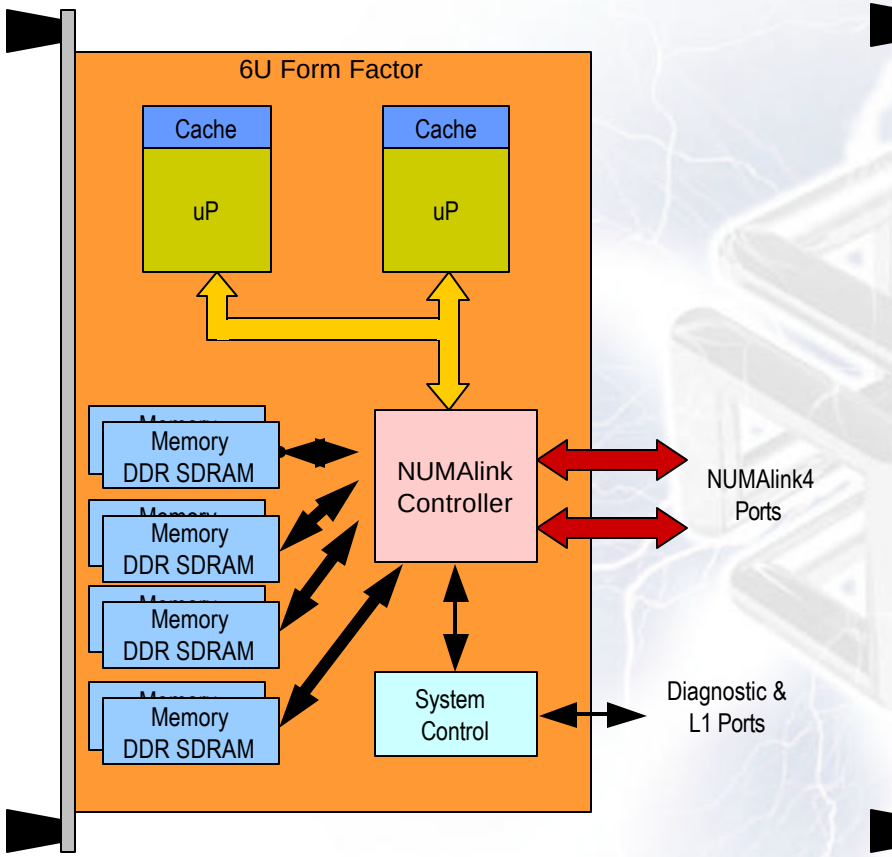


**Scalable General Purpose Processor/Memory Blade**

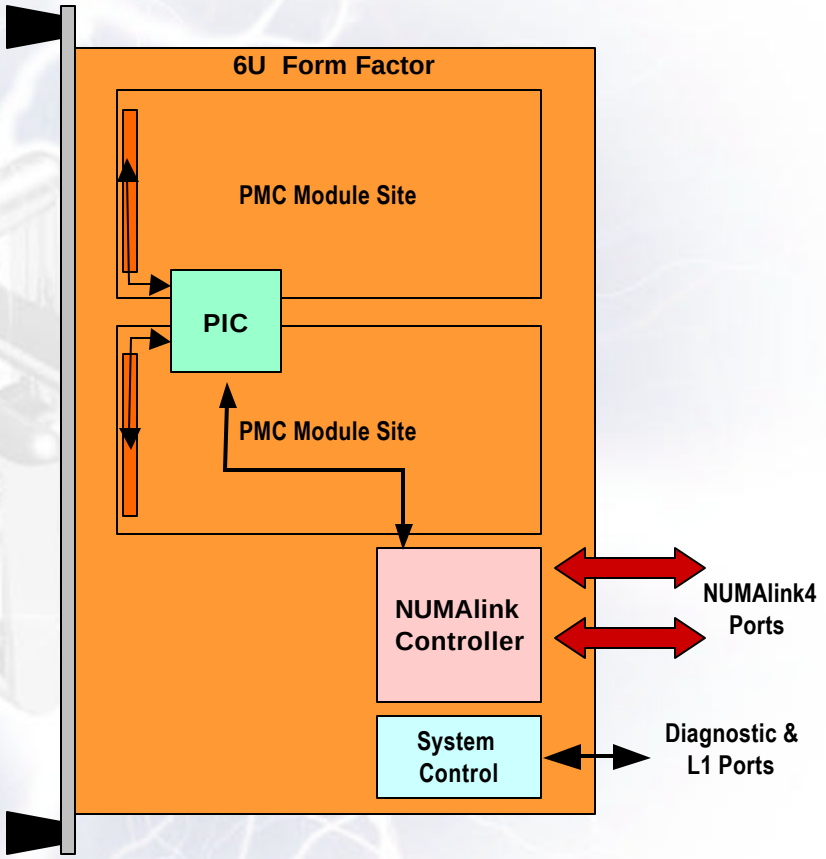


**General Purpose I/O Blade**

Igniting Innovation and Leadership

# Embedded High Performance Computing
## Family of 6U Form Factor Modules

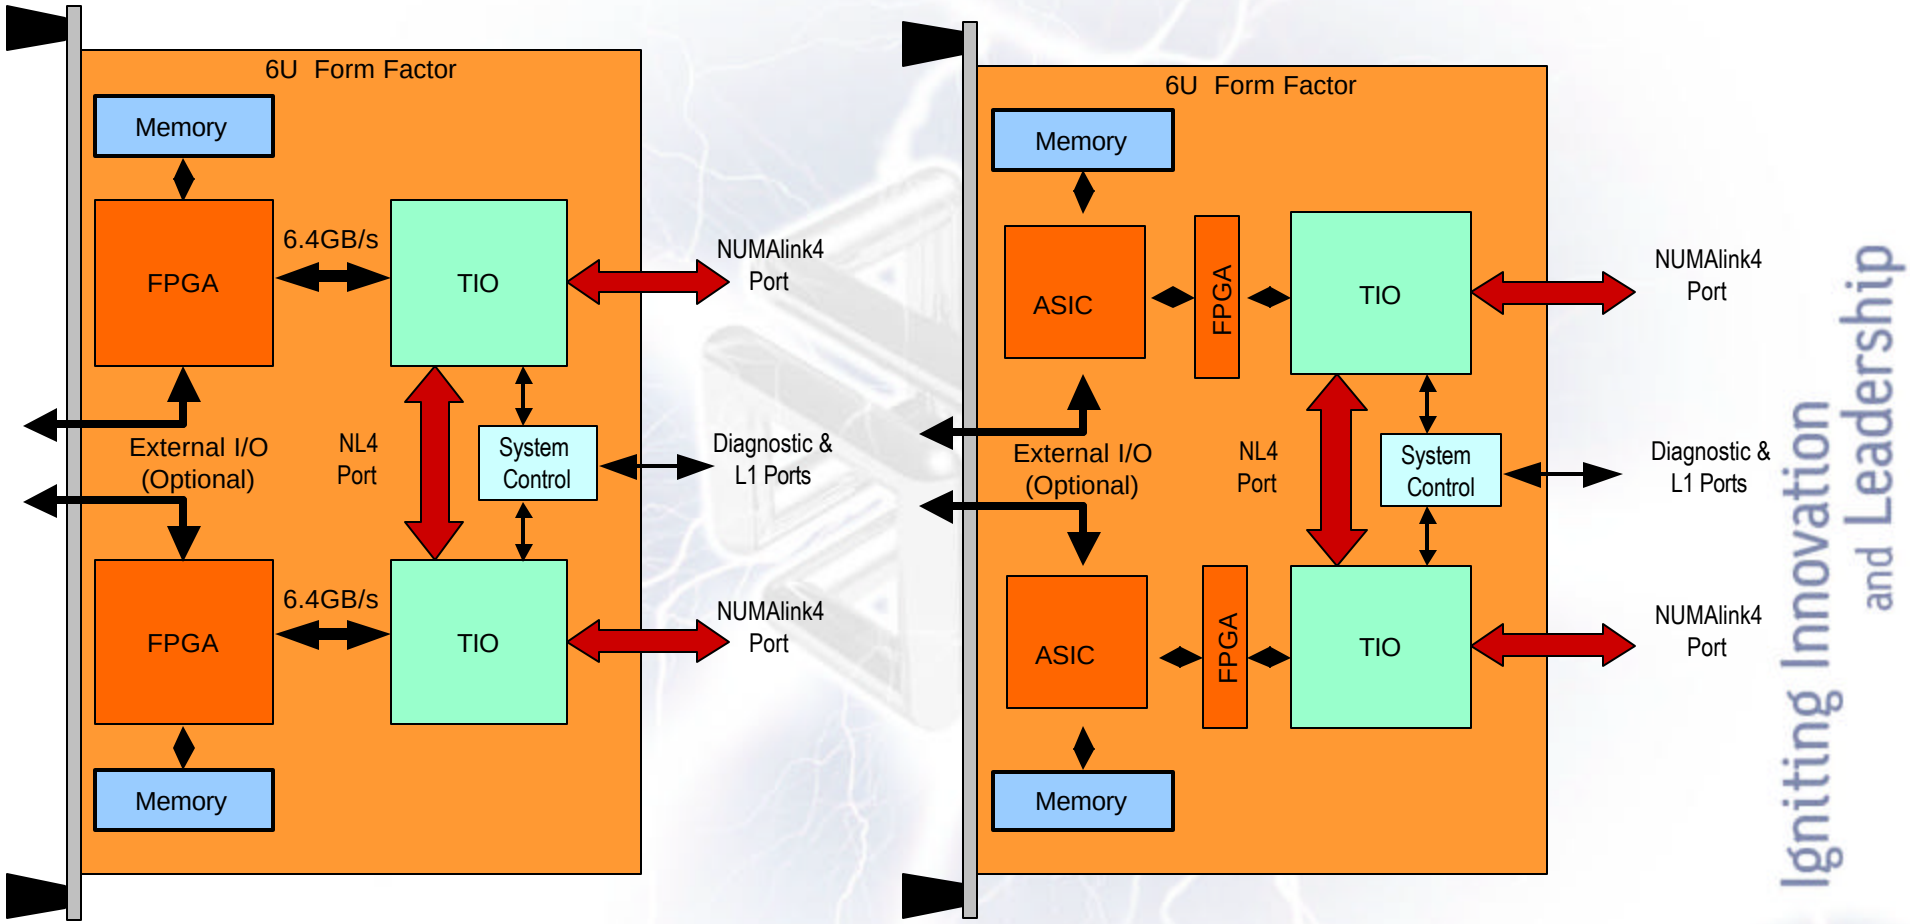


**Mission-Specific Accelerator and/or I/O Blades**

Igniting Innovation and Leadership