



# HP-SEE



HP-SEE receives EC support through FP7  
under the "Research Infrastructures" action

**HPSEE TRAINING**

**30 NOVEMBER 2010, SOFIA, IICT-BAS**

**SPONSORED BY:**

**sgi**®

scripto

enterprise content management  
information security  
business process management



# **Structures of Large HPC Centers in Europe**

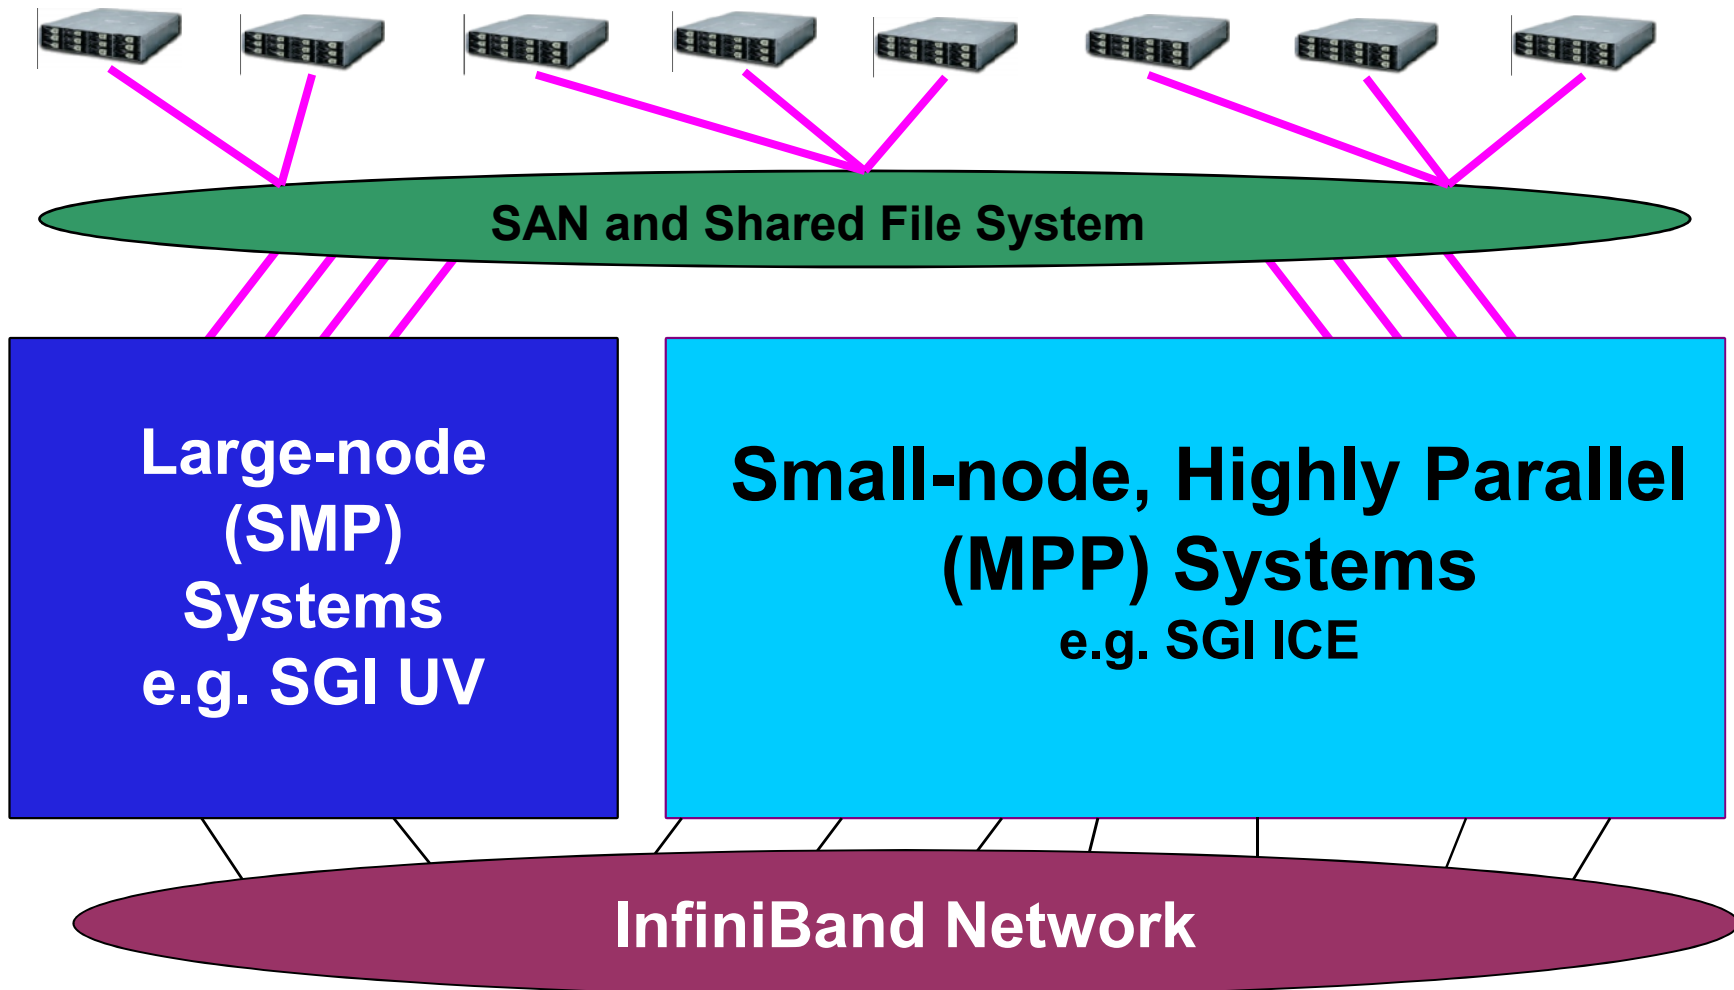
**Robert Uebelmesser**  
([rue@sgi.com](mailto:rue@sgi.com))



# Topics

- **SGI Customers, Target Markets, and Product Focus**
- **A Key Challenge for Future HPC Systems – The System Interconnect**
  - Applications Analysis
  - The SGI ICE Approach
  - The SGI UV Approach (solves the large memory problem at the same time)
- **A Futuristic Data Center – SGI ICE Cube**
- **The Structure of large HPC Data Centers in Europe**

# Very Large HPC Centers in Europe





## Relevant European SGI Sites in the Top500 list

# SGI ICE at Genci/CINES in France





# Altix<sup>®</sup> ICE:

## CINES Grand Equipement National de Calcul Intensif (GENCI)



- **SGI Altix ICE with**
  - 12,288 processor cores ,
  - 50 TB of memory
  - 500TB SGI<sup>®</sup> InfiniteStorage 4600 solution
  - Capable of 147 Tflops

- **This year upgraded with an additional**
  - 10752 processor cores
  - 200 Tbytes of storage

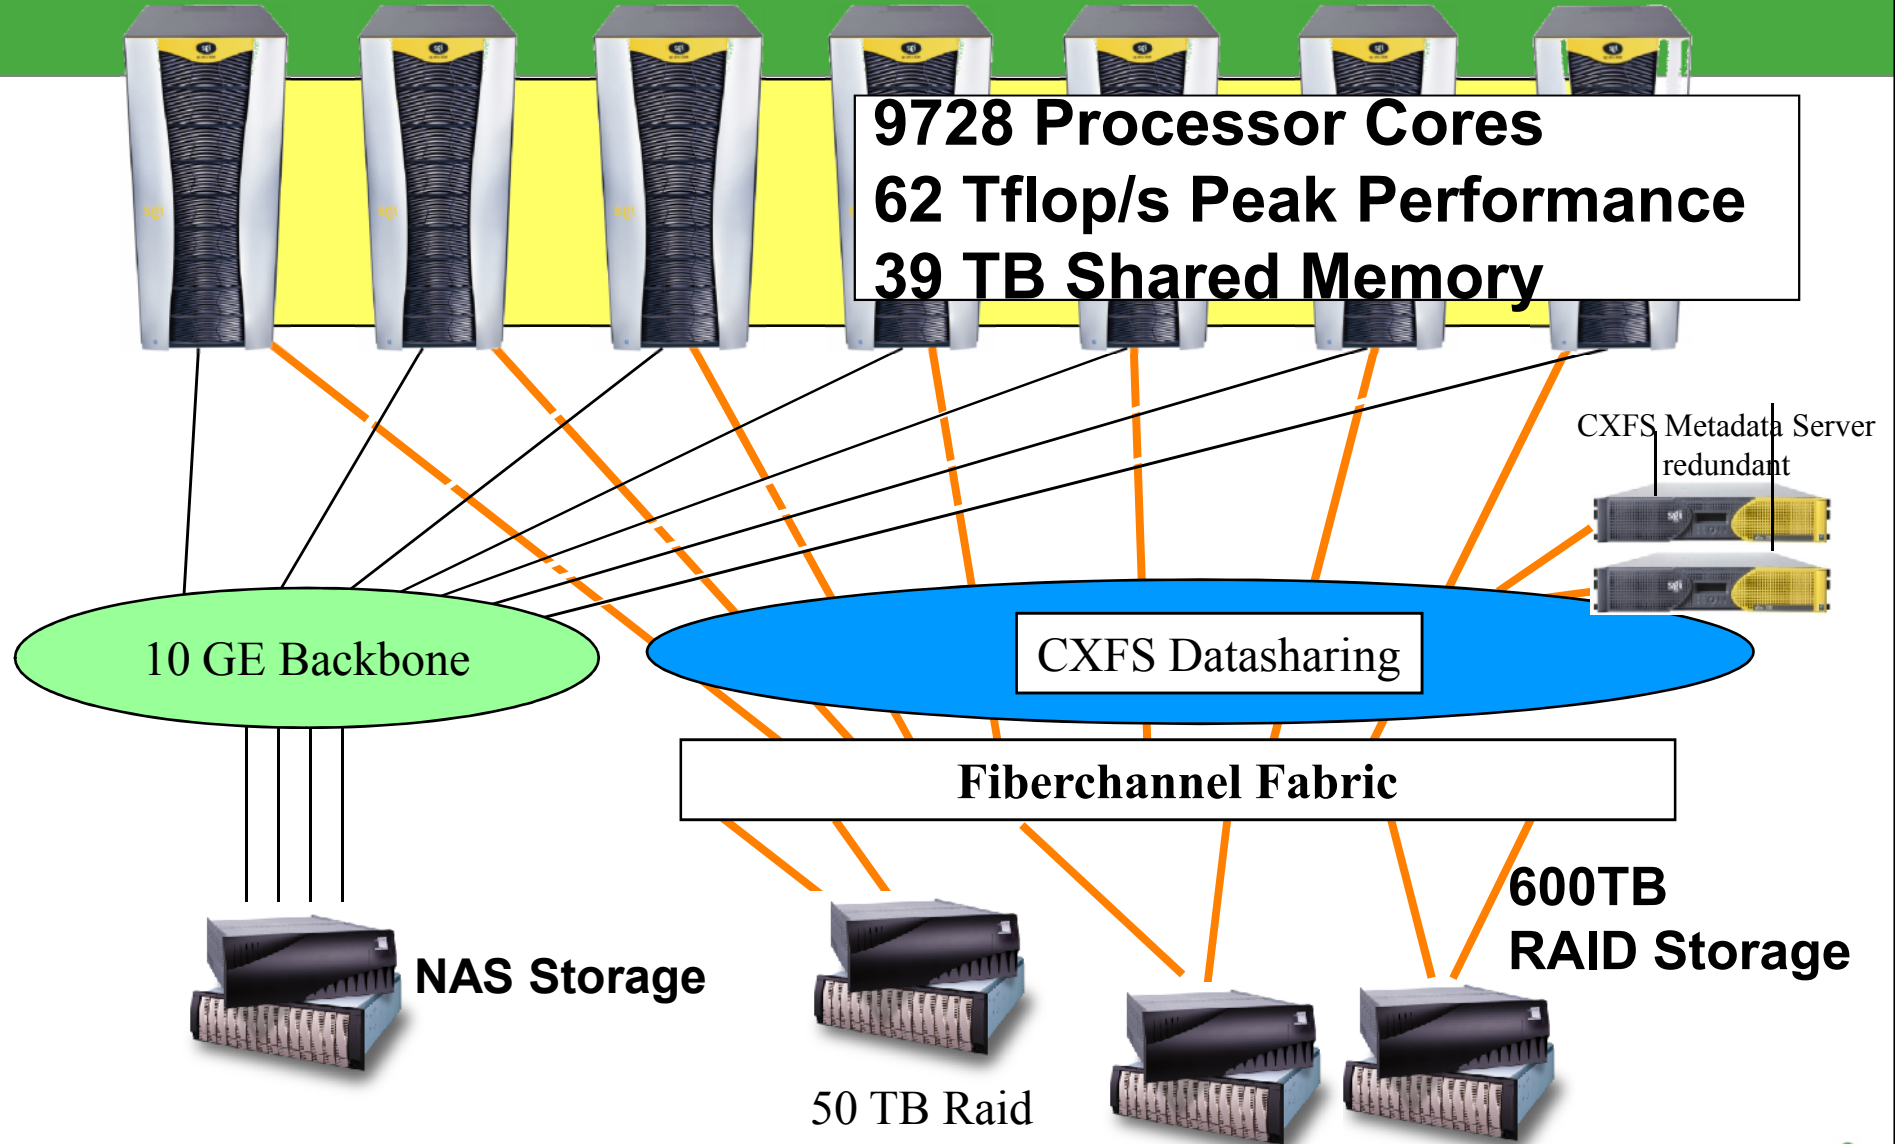


# SGI Altix 4700 System at LRZ (Munich)

## The World's Largest Shared Memory System



# SGI Altix 4700 Configuration at LRZ



# Altix<sup>®</sup> ICE: TOTAL Oil Company

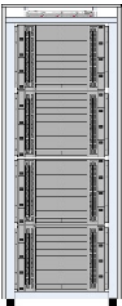
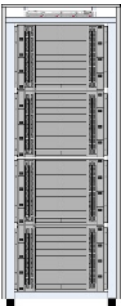
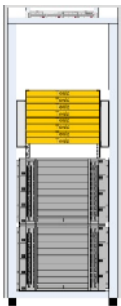


- **SGI<sup>®</sup> Altix<sup>®</sup> ICE 8200 EX with**
  - 10240 processor cores
  - 20 Terabytes (TB) of main memory
  - 1.5 Petabytes of SGI InfiniteStorage
- # 10 on TOP 500 at installation time
- #1 industrial system in TOP 500 at installation time

- **this year installed an upgrade with an additional**
  - 6656 processor cores
  - 256 GPGPUs
  - 2.3 Pbytes of storage

# SGI Configuration at Total

- Upgrade with CPUs and GPUs
- One single IB Fabric with existing racks

Existing		Extension 2009			
20 Racks ICE Harpertown		12 Racks ICE Nehalem		8 Racks ICE Hybrid	
	<ul style="list-style-type: none"> <li>- 2560 processors Harpertown 3GHz</li> <li>- 10240 cores</li> <li>- 20 TB de mem</li> <li>- 2 GB/core</li> <li>- 1280 blades Colfax-S</li> </ul>		<ul style="list-style-type: none"> <li>- 1536 processors Nehalem 2.8 GHz</li> <li>- 6144 cores</li> <li>- 18 TB de Mem</li> <li>- 3 GB/core</li> <li>- 768 blades Bloomer-S</li> </ul>		<ul style="list-style-type: none"> <li>- 256 processors Harpertown 3GHz</li> <li>- 1024 cores</li> <li>- 2 TB de mem</li> <li>- 2 GB/core</li> <li>- 128 blades Colfax-D</li> <li>- 64 <b>Tesla</b> S1070-500</li> <li>- 256 GP-GPU</li> </ul>



# HLRN: North German Computing Alliance

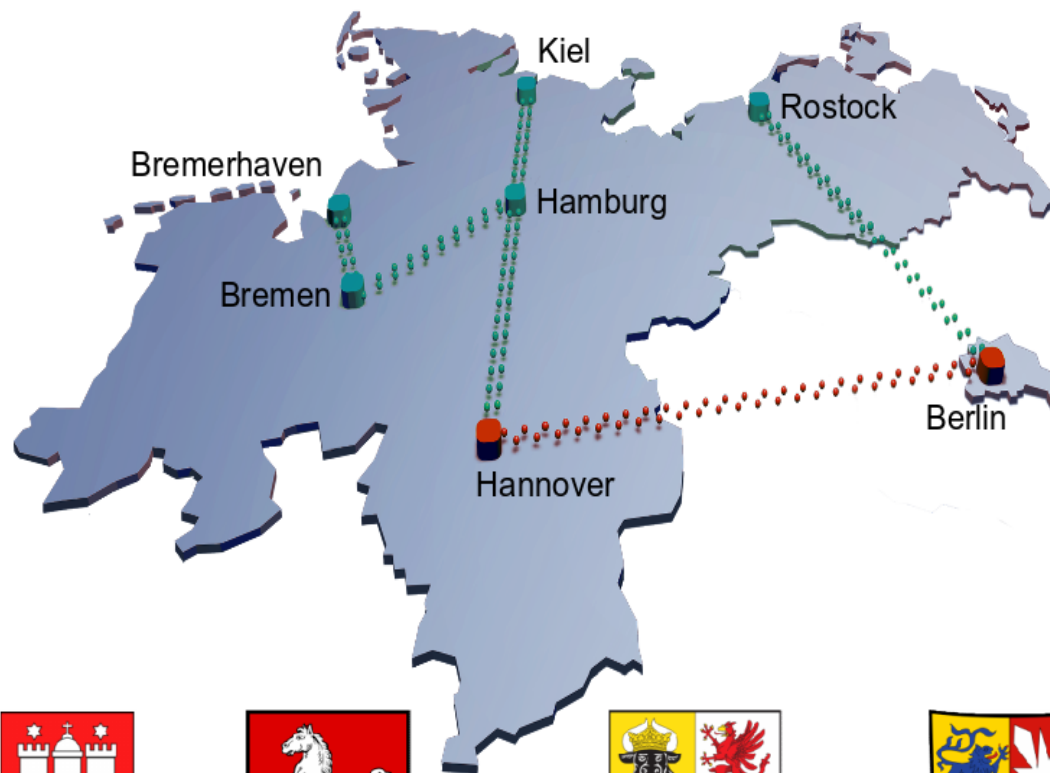


## The Solution

- Hybrid computing solution with SGI® Altix® ICE and SGI® Altix® XE cluster
- 25,000 processor cores, 90 Terabytes (TB) of main memory and 2.3 Petabytes of SGI InfiniteStorage storage capacity
- over 300 Teraflops peak performance



## Norddeutscher Verbund für Hoch- und Höchstleistungsrechnen




# HLRN

“North German Supercomputing Alliance”

Berlin

Hanover

**One Grid – One Workload – One Filesystem**



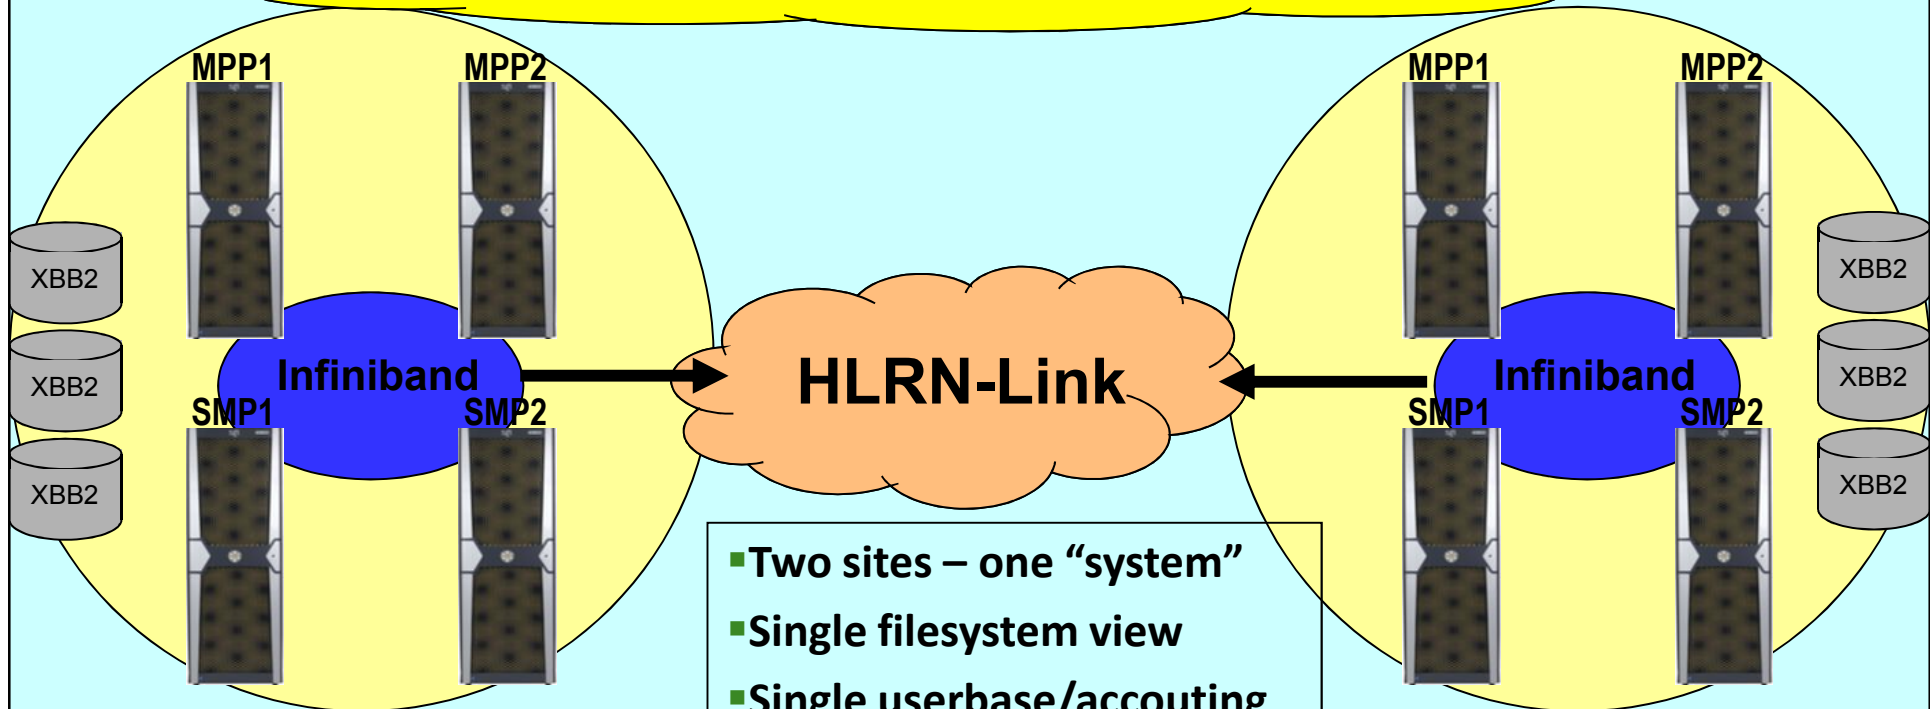


# HLRN2 Grid – Single System View

Berlin

Hannover

**One Grid – One Workload – One Filesystem**



- Two sites – one “system”
- Single filesystem view
- Single userbase/accouting
- Single workload managed across two sites



- **Total Workflow Solutions for**
  - **The Roshydromet Main Computer Center Moscow**
  - **The Roshydromet Novosibirsk Regional Center**
  - **The Roshydromet Khabarovsk Regional Center**

# The Roshydromet Main Computer Center Moscow

- **Altix 4700 with**

- 1664 cores of 1.66 GHz Montvale
- 6.6 TBytes Globally Shared Memory



- **SGI ICE 8200 with**

- 1416 cores of 2.8 GHz Harpertown

- **Storage**

- Infinite Storage 4000: 60 Tbytes
- Infinite Storage 10000 150 Tbytes
- ADIC scalar i500 86 Tbytes

- 12 Altix 450 functional servers

- 8 Altix XE240 functional servers

- 9 Altix XE210 functional servers

- CXFS, DMF, Cluster Manager, Networker, PBSPro, Scal

# The Roshydromet Regional Computer Centers

## Novosibirsk, Khabarovsk, St. Petersburg

- **Altix 4700 with**

- 104 cores of 1.6 GHz Montecito
- 208 GB Globally Shared Memory

- 5 Altix 450 functional servers
- 4 Altix XE240 functional servers
- 10 Altix XE210 functional servers

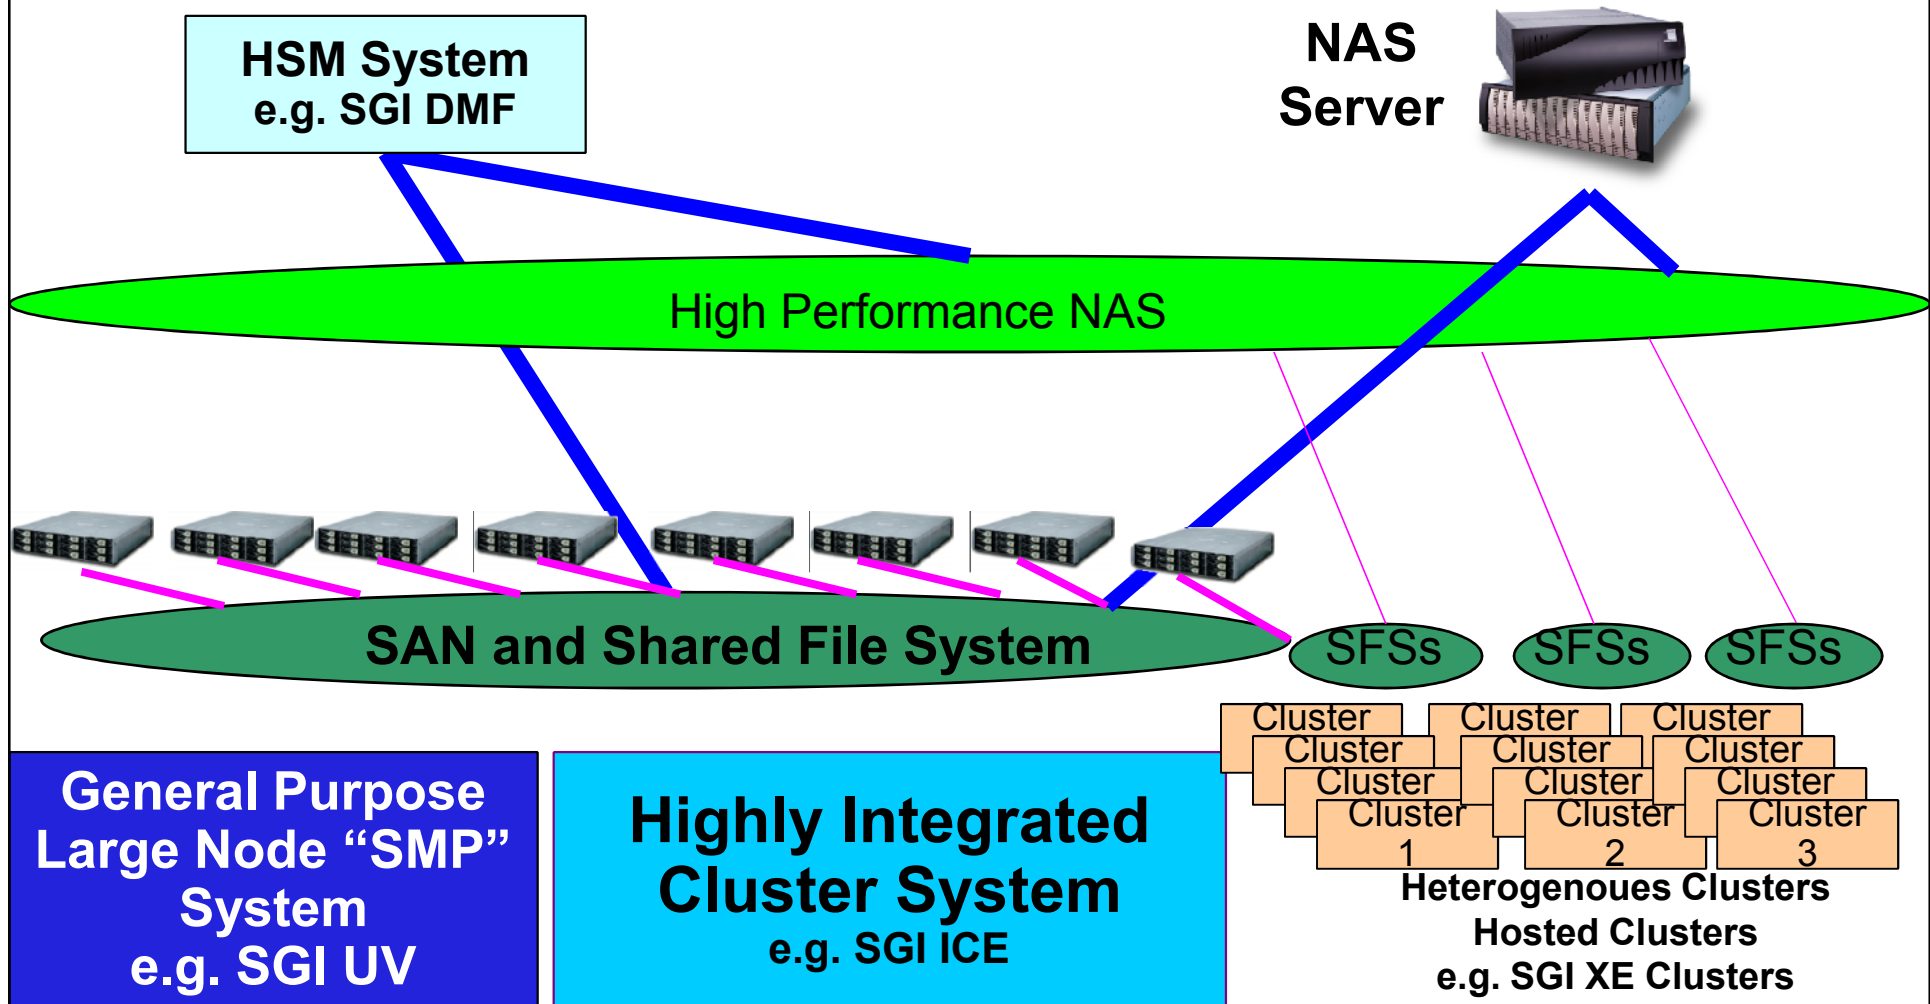
- **Storage**

- Infinite Storage 4000f: 12 Tbytes
- ADIC scalar i500 32 Tbytes

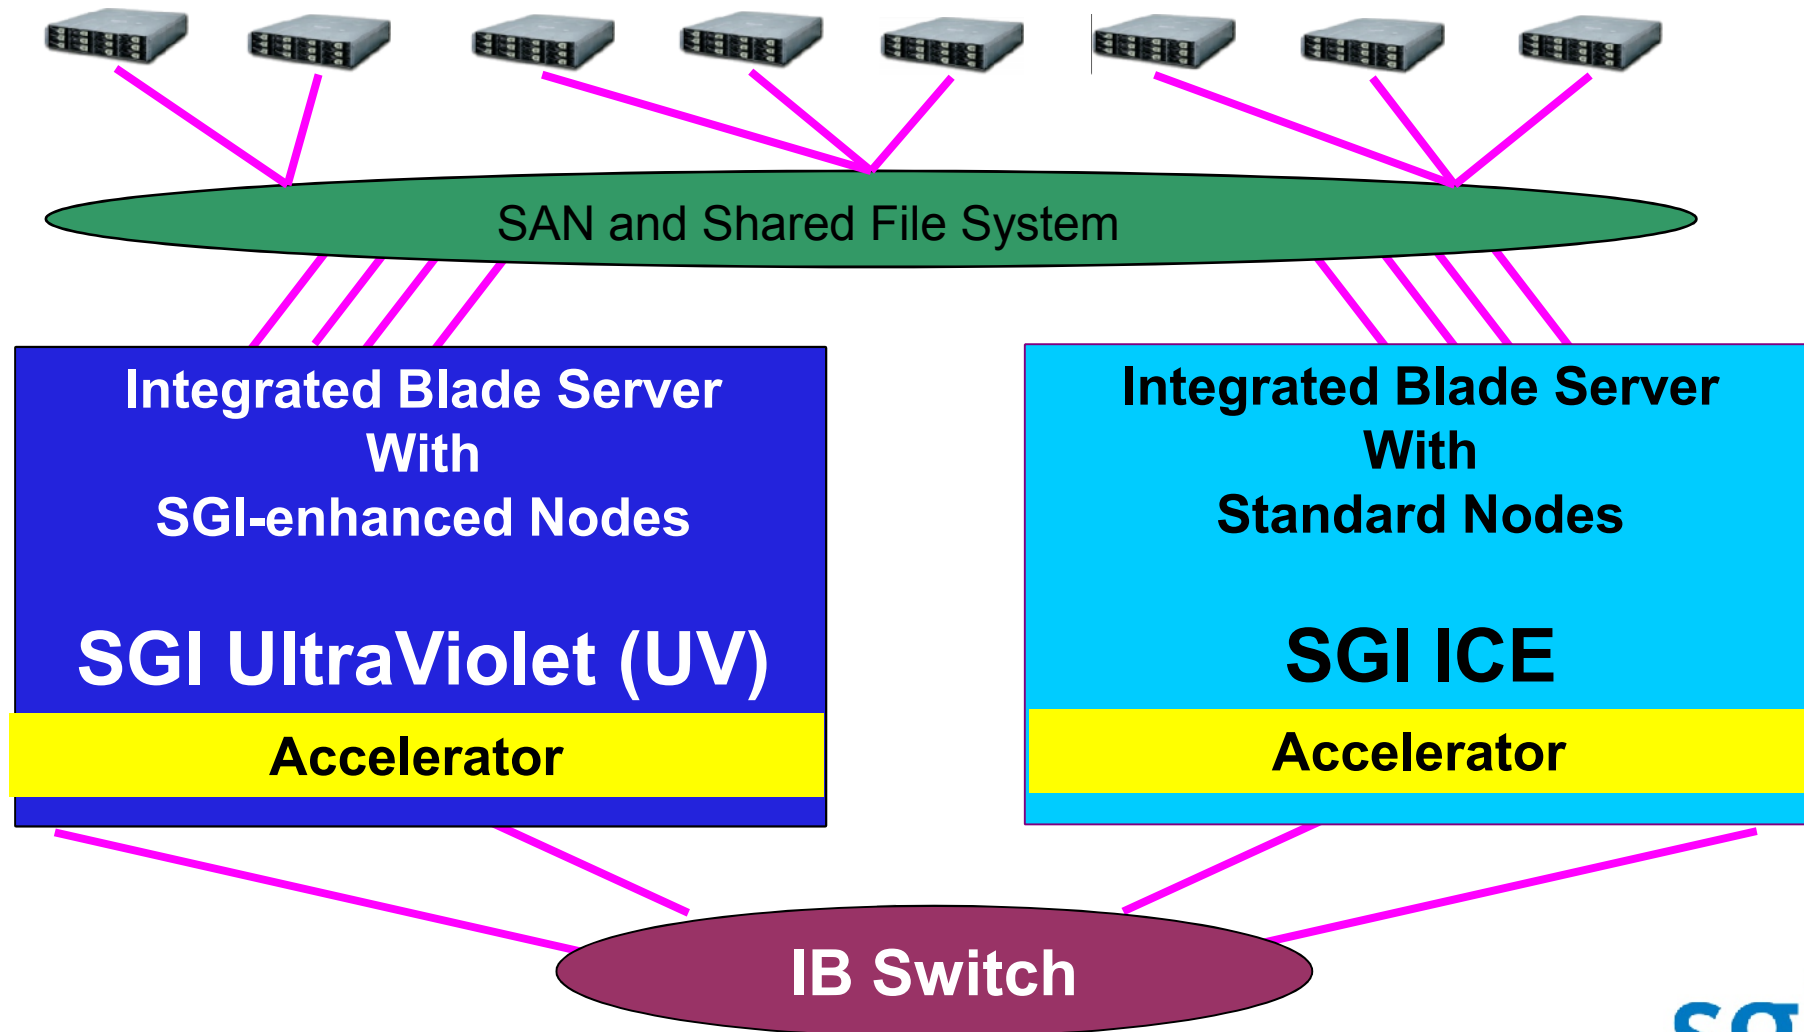
- **CXFS, DMF, Cluster Manager, Networker, PBSPro, Scali**



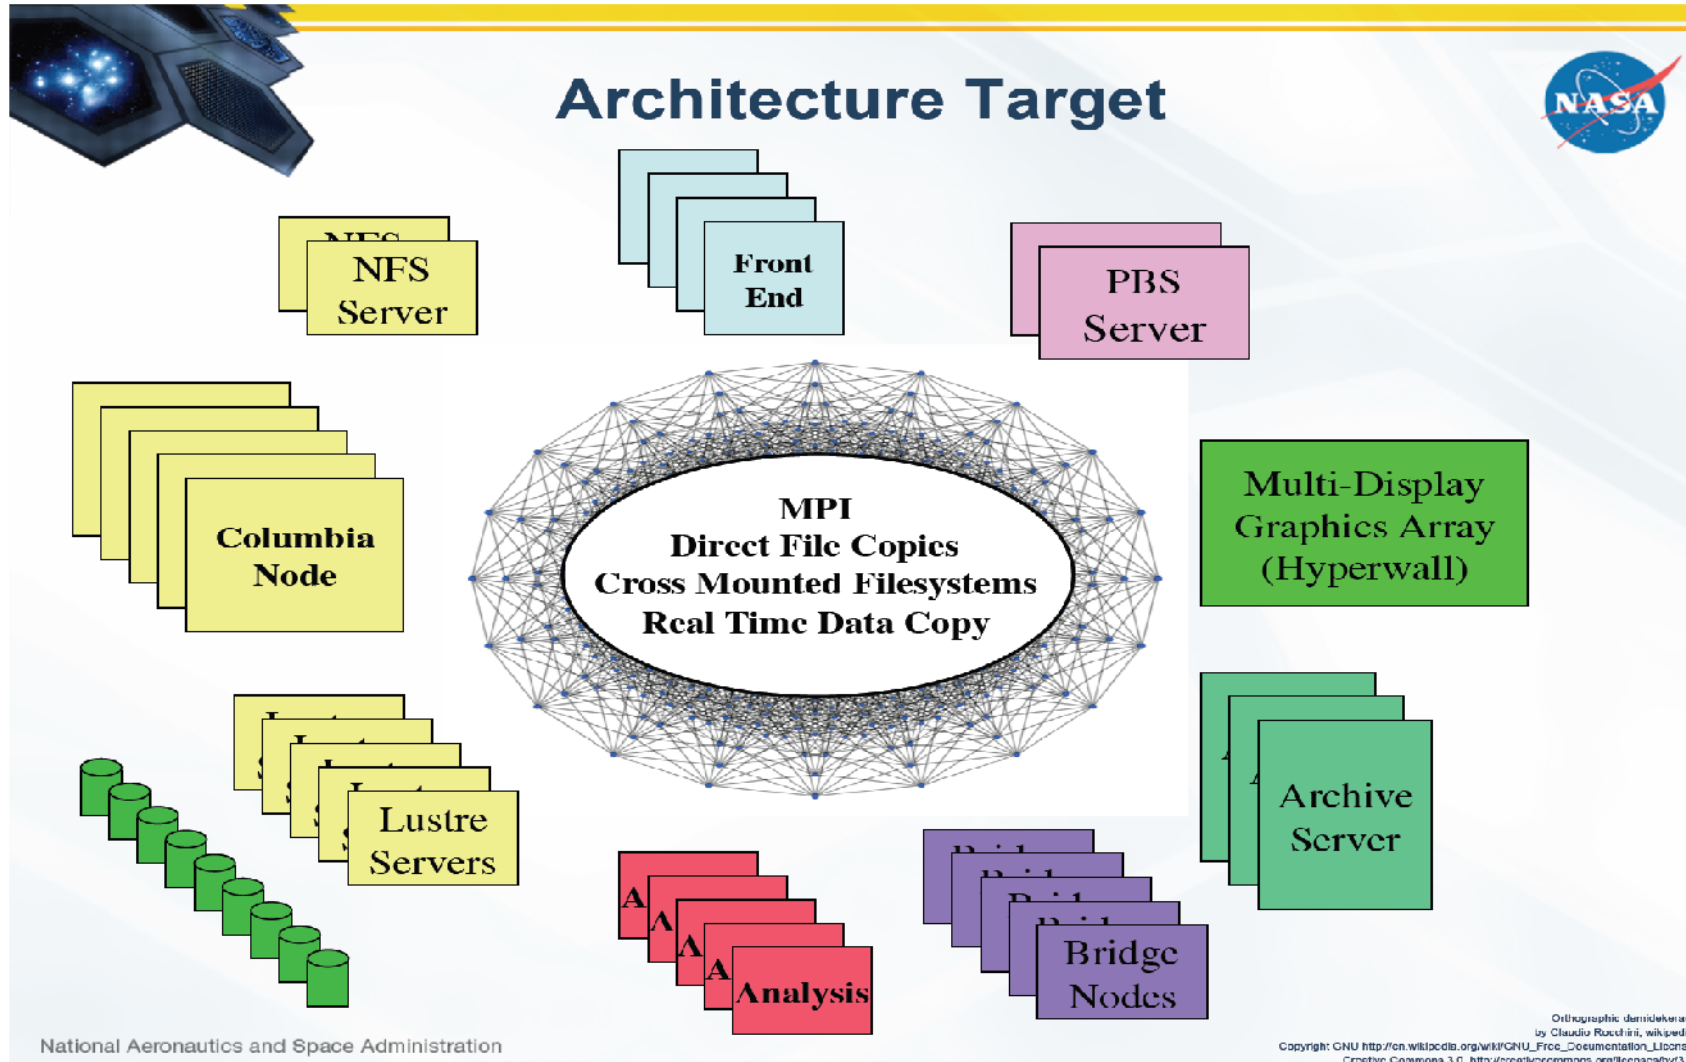
# Very Large HPC Centers in Research



# LRZ/CINES PRACE Project with SGI



# SGI System Complex at NASA/Ames





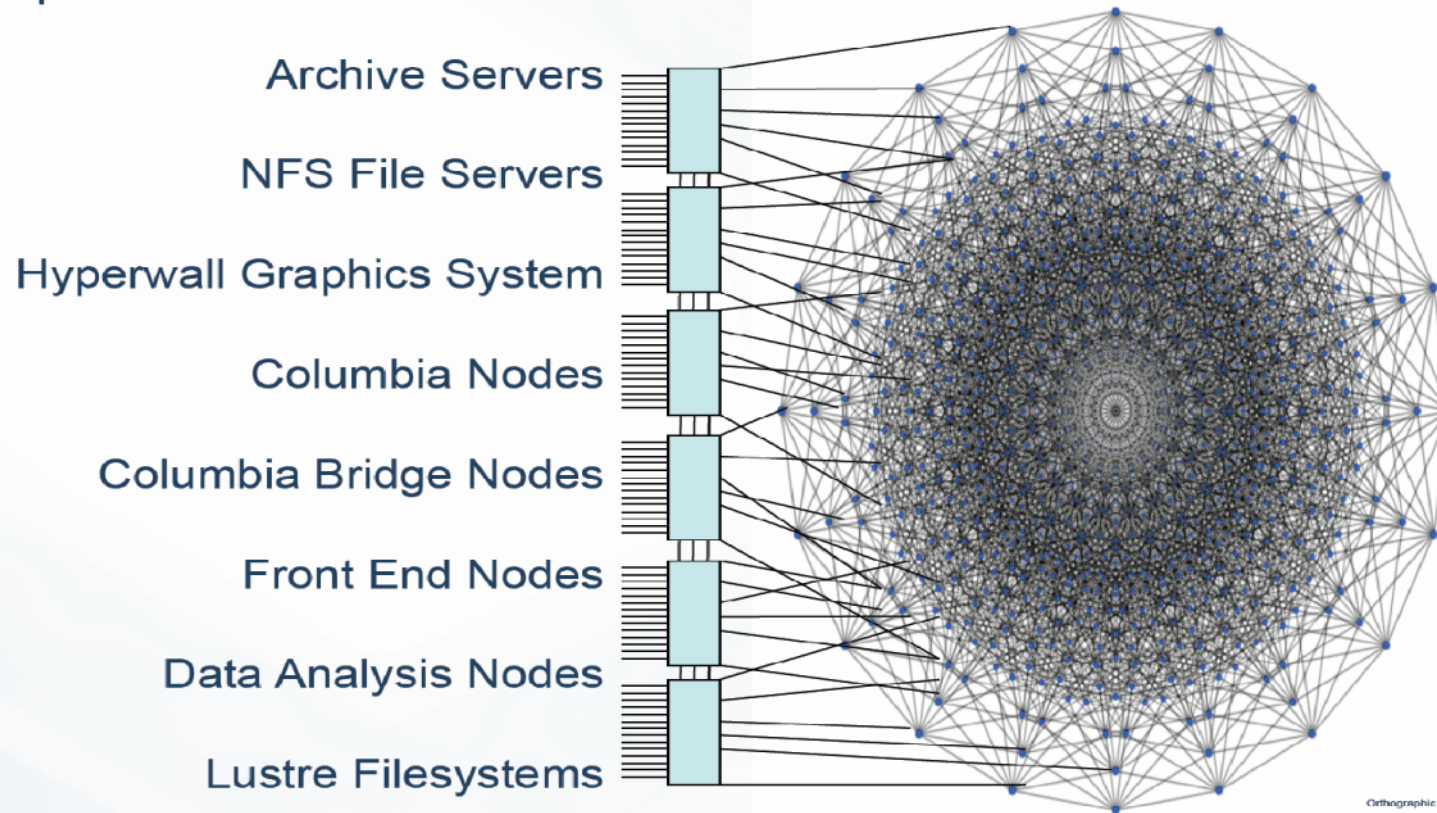
# SGI System Interconnect at NASA/Ames



## Infiniband Subnet LAN



LAN Implemented with out board IB switches



Orthographic dimidiate  
by Claudio Rocchini, wikipedia





## NAS September 2008



- Pleiades (608TF) (44934 SBU)
  - 6400 dual-socket blades
  - 12,800 quad-core Intel Xeon processors
  - 51,200 cores
- Columbia (86TF) (13716 SBU)
  - 22-node SGI Altix Supercluster
    - 10 512c 3700s
    - 8 512c 3700BX2s
    - 1 512c 4700
    - 2 1024c 4700s
    - 1 2048c 4700
- Schirra (4.8TF) (1016 SBU)
  - 40 IBM 9118-575 nodes
  - 320 dual-core IBM processors

**699 TF**  
**59666 SBUs**

# Summary

- **SGI Customers, Target Markets, and Product Focus**
- **A Key Challenge for Future HPC Systems – The System Interconnect**
  - Applications Analysis
  - The SGI ICE Approach
  - The SGI UV Approach (solves the large memory problem at the same time)
- **A Futuristic Data Center – SGI ICE Cube**
- **The Structure of large HPC Data Centers in Europe**

