



# Networking Update

George Hyman  
gmh@sgi.com

Network Engineering Director, SGI

SUMMIT 2001

43rd CUG Conference on High Performance Computing & Visualization



# SGI Networking Focus

---



- Focus on high performance local area network technology and clustering
- Partner with 3rd parties for commodity PCI hardware
- Bring clustered solutions to market through Bandwidth Solutions of appropriate switching technology<sup>2</sup>



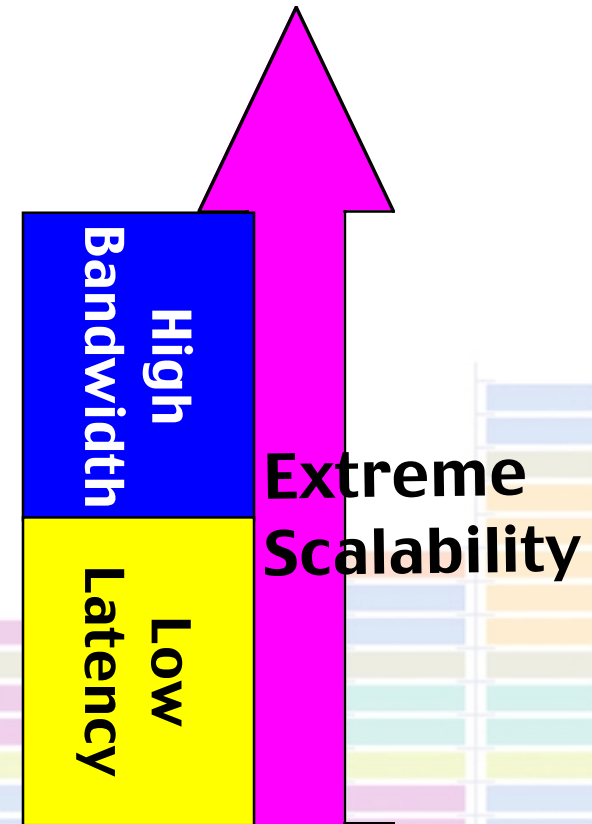
SUMMIT 2001

43rd CUG Conference on High Performance Computing & Visualization

# Best-in-Class Performance



**Pioneer in Intelligent  
Networking Technology**  
**Industry-Leading, Award-  
Winning Performance**  
**Scalable Solutions from Small  
Workgroup Servers through  
Super Computers**



SUMMIT 2001

43rd CUG Conference on High Performance Computing & Visualization

# Connectivity Comparison



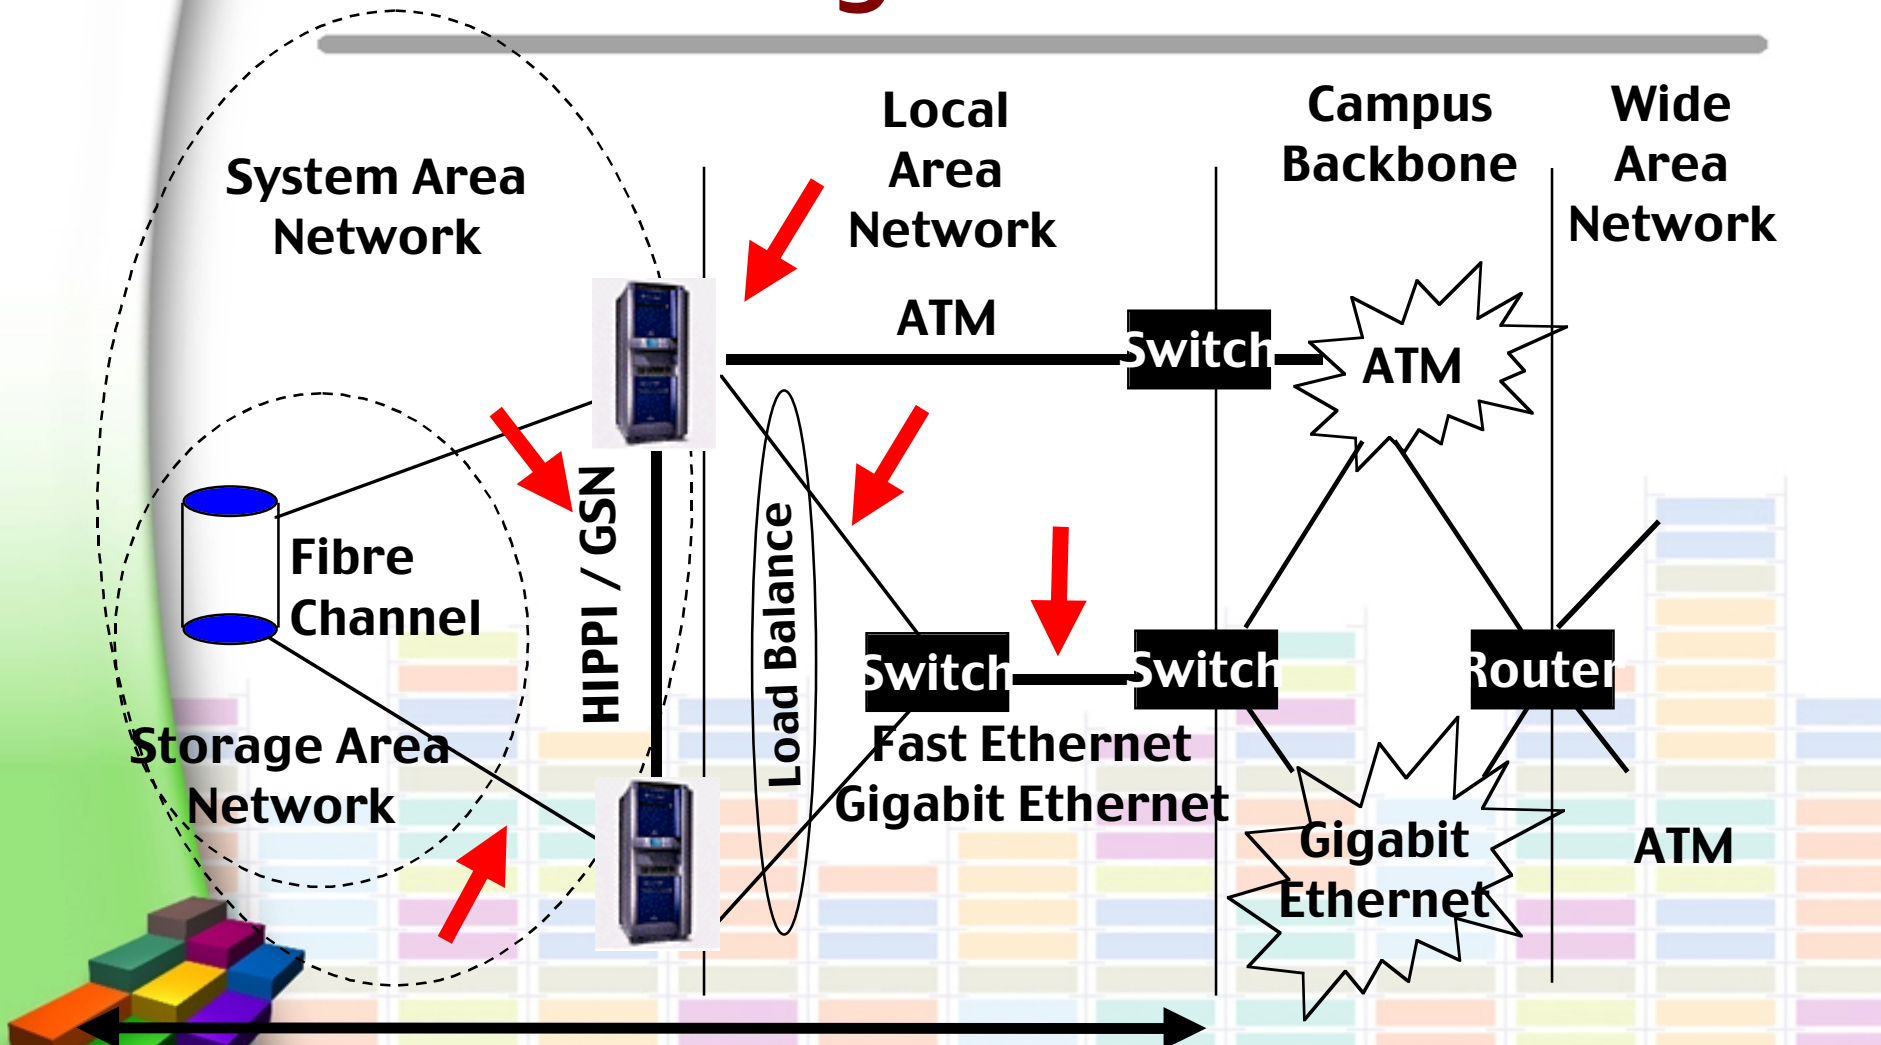
<i>Technology</i>	<i>Bandwidth (Mbps)</i>	<i>Latency (<math>\mu</math>s)</i>	<i>CPU util</i>
<i>GSN</i>	<i>6400</i>	<i>6 <math>\mu</math>s</i>	<i>&lt;10%</i>
<i>Myrinet</i>	<i>2000</i>	<i>10 <math>\mu</math>s</i>	<i>&lt;10%</i>
<i>GigE</i>	<i>1000</i>	<i>200 <math>\mu</math>s</i>	<i>125%</i>
<i>Fast Ether</i>	<i>100</i>	<i>200 <math>\mu</math>s</i>	<i>low</i>



SUMMIT 2001

43rd CUG Conference on High Performance Computing & Visualization

# Where Does SGI Networking Excel?



SGI Networking Solutions SUMMIT 2001

43rd CUG Conference on High Performance Computing & Visualization

# ***Gigabyte System Network™, GSN™***

**sgi™**

- **Is the highest bandwidth and lowest latency interconnect standard**
  - Provides full duplex dual, 6400 megabit (800 megabyte) per second of error-free, flow controlled bulk data
  - Under 10μs round-trip latency [using ST & OS bypass]
- Provides for interoperability with other networking media through standard both IP and ST standard protocols.



**SUMMIT 2001**

**43rd CUG Conference on High Performance Computing & Visualization**



# GSN Updates

- GSN 2.1 – supports Origin 3000, Onyx 3
- GSN 3.0 – libst productized, continued performance work
- ST protocol being reviewed/updated for Infiniband compatibility
- Latest performance results:
  - ST – 775 megabytes/sec using 6% CPU!***
  - TCP – 507 megabytes/sec using 50% CPU***

SUMMIT 2001

43rd CUG Conference on High Performance Computing & Visualization



# GSN – Where to use?

---

The technology is ideal wherever organizations require timely movement of large amounts of information including:

- scientific and technical computing
- HDTV
- data mining
- transaction processing
- video and film archiving
- storage management

SUMMIT 2001

43rd CUG Conference on High Performance Computing & Visualization





# Gigabit Ethernet



## *Market Applications*

- **Backbones, Server Interconnects**
  - **Horsepower to fulfill LAN bandwidth requirements**
- **High-Performance Workgroups**
  - **Greater deployment w/higher density GbE switches**
- **Clusters**
  - **Point to Point system interconnects, small gigabit switches**
- **Pinpoint 10/100 Ethernet bottlenecks**
  - **Key Target Markets are File Serving and Backup Environments**
  - **Gigabit Ethernet is Ethernet...only faster!**



SUMMIT 2001

43rd CUG Conference on High Performance Computing & Visualization

# Gigabit Ethernet Update

---



- 6.5.12 has 2 significant bug fixes – patches available back to 6.5.9
- Performance work going on to reduce CPU loading
- Both Copper and Fiber available
- Working on new baseIO to include Gigabit Ethernet



SUMMIT 2001

43rd CUG Conference on High Performance Computing & Visualization

# Myrinet

---

- OEM from Myricom, Inc.
- Supported for cluster interconnect
- Hardware and software sold by SGI
- Included in Origin 3200C
- Supported for cluster interconnect only on Origin 3000/Onyx 3



SUMMIT 2001

43rd CUG Conference on High Performance Computing & Visualization

# Network Load Balancing Software (NLBS)



## *Performance Boost for Existing Ethernet based Networks*

- **Dynamic Load Distribution across multiple Ethernet interfaces**
- **Balances both inbound and outbound traffic**
- **Eliminates Single-Point of Failure at NIC**
- **No configuration for L2 devices**
- **Interoperability with heterogeneous clients**
- **Supports 2-40 adapter ports per subnet**
- **Available for all Origin/Onyx/Challenge**
- **Evaluating 802.3ad Link Aggregation**



SUMMIT 2001

43rd CUG Conference on High Performance Computing & Visualization

# Cluster Interconnects

---

- GSN – fastest, lowest latency. XIO only  
STP or TCP/IP
- Myrinet – lower cost, GM protocol
- Gigabit Ethernet – TCP/IP or STP
- NUMAflex – MPI and IP between partitions



SUMMIT 2001

43rd CUG Conference on High Performance Computing & Visualization

# Networking Hardware Roadmap

---



- LAN consolidation to Ethernet, mostly Gigabit
- 10 Gigabit Ethernet as backbone or switch to switch in short-medium term
- Infiniband as computer room and cluster interconnect.
- GSN as a precursor to Infiniband using STP
- Investigating IP stack on chip/board



SUMMIT 2001

43rd CUG Conference on High Performance Computing & Visualization



# Networking Software Roadmap

---



- IPv6 – tested at Connectathon, Beta soon
- IPsec
- NFSv4
- Continue RFC compliance work
- Continued work with the Linux community



SUMMIT 2001

43rd CUG Conference on High Performance Computing & Visualization