NASA Optical Networking Research

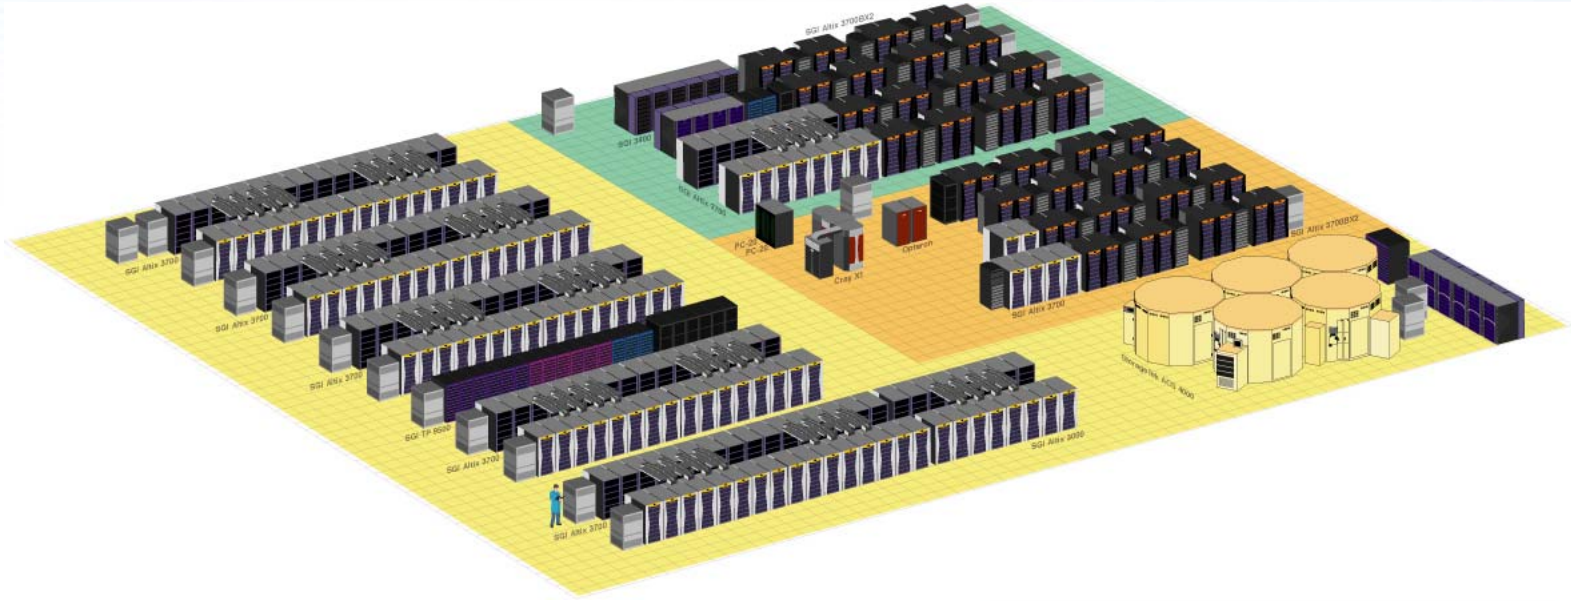

Kevin L. Jones

NASA Advanced Supercomputing Division

NASA Ames Research Center

8 September 2006

NASA Research and Engineering Network (NREN) Overview

- NREN is a research network, that is currently focused on supporting NASA's High-End Computing (HEC) networking requirements
- Specifically, the Columbia Supercomputer is located at NASA Ames Research Center in Mountain View, CA
- NASA Advanced Supercomputing (NAS) Division manages both NREN and the Columbia Supercomputer

NREN Target CY06

10G waves at the core, dark fiber to end sites

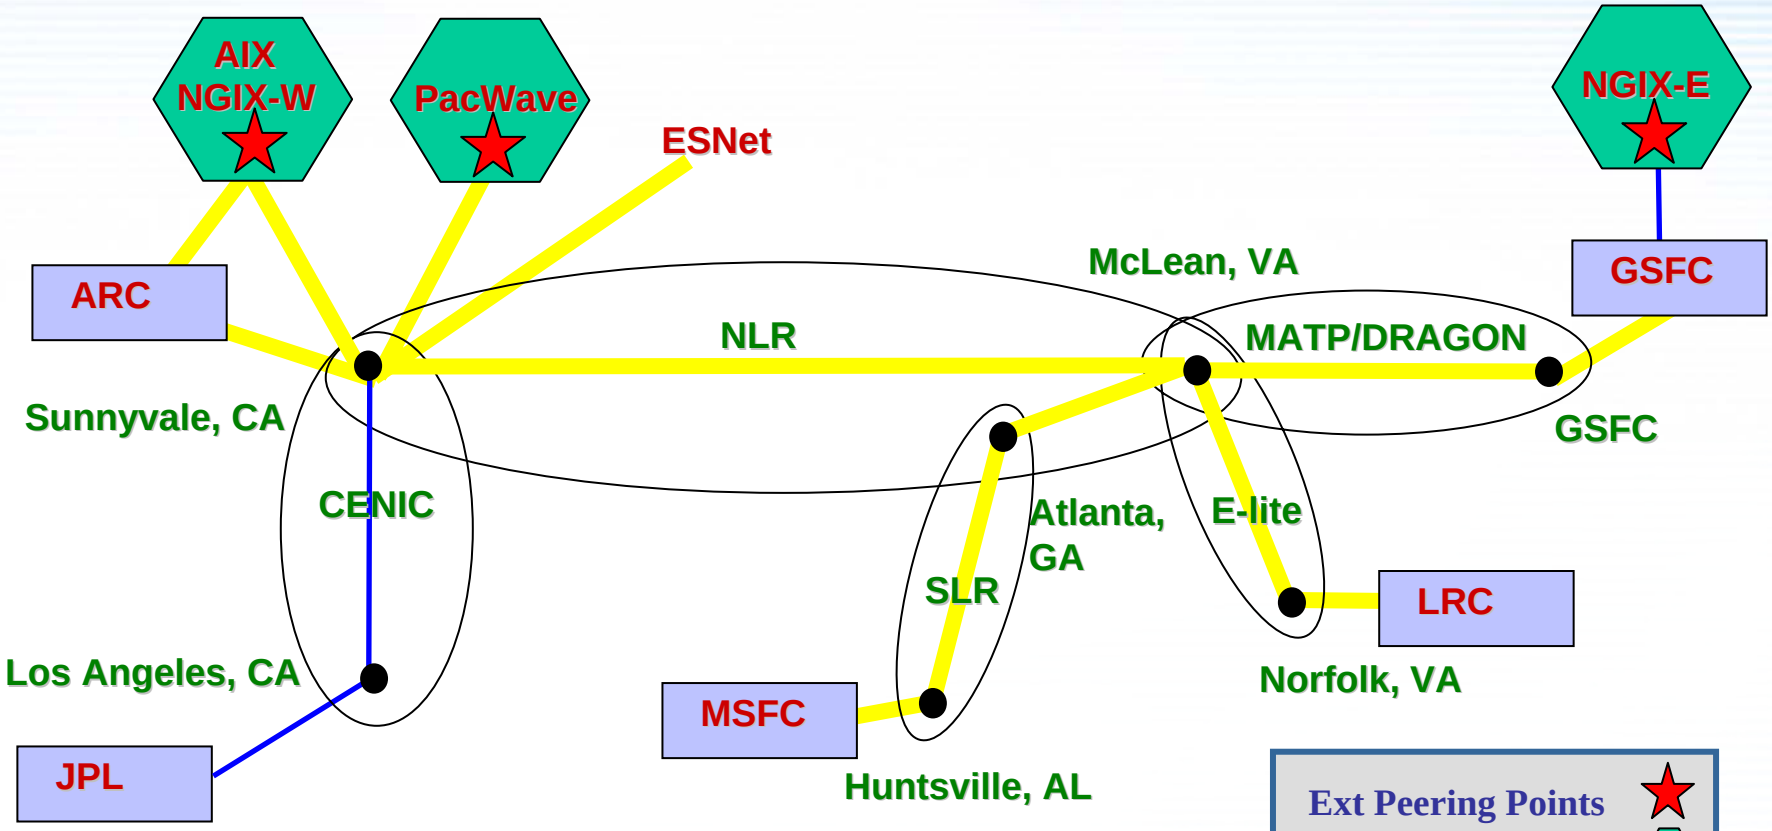

- National and Regional optical networks provide links over which 10 Gbps and 1 Gbps waves can be established.
- Distributed exchange points provide interconnect in metro and regional areas to other networks and research facilities

International Peering

- NREN established presence in Sunnyvale, CA recently and now is a member of the Pacific Wave distributed international peering exchange
- Networks interested in setting up bilateral peering should email: peering@nren.nasa.gov
- International Peering
  - Pacific Wave major source
  - Transit over another peer network

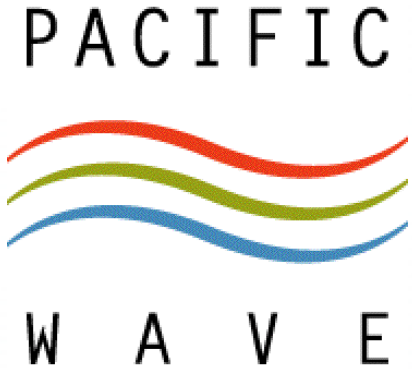

Columbia Supercomputer

- SGI Altix 3700 Architecture
- ~1 PB of online disk space
- 10,240 1.5-1.6 GHz CPUs configured as twenty 512 CPUs
- Each 512 CPU node:
  - 1 TB of shared memory
  - Five x 1 GigE
  - 1 or 2 10 GigE
  - Multiple Fiber Channel links
  - Multiple Infiniband connections
- 2048 CPU subcluster links 4 nodes via NUMALink4

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

RDMA Research

- NREN is interested in developing methodologies for increasing bandwidth for NASA applications

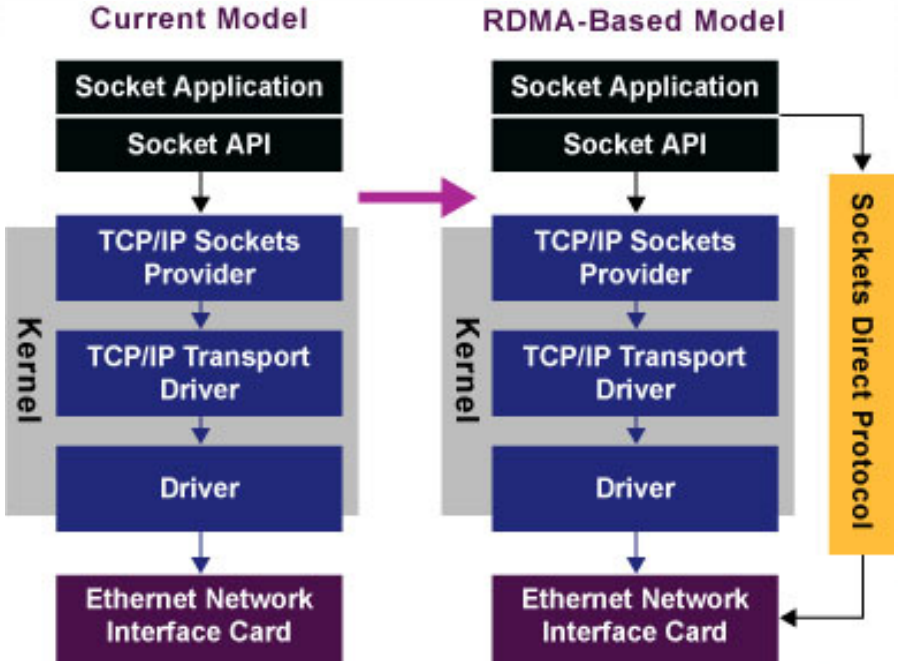

- Remote Direct Memory Access (RDMA) allows direct access to a processing node's hardware such as memory and disk systems, bypassing the TCP/IP stack

Infiniband Research

- Infiniband (IB) has become a well established technology, in the clustering world, to achieve RDMA hardware acceleration
- Selected for transport because:
  - Mature RDMA support on both hardware and software stack
  - Broad acceptance in supercomputing applications as a cluster and/or storage interconnect
  - Offers low-latency (< 5 μs) end-to-end transport locally
  - Cost savings from integrating clustering, storage, and WAN I/O onto a single network

Infiniband Research

Motivation:

- Distributed application processing between non-local nodes
- Direct memory sharing - supercomputer memory shared directly with visualization memory without going through the TCP/IP stack

Approach:

- Connect two Columbia nodes via Infiniband
- Used Obsidian's Longbow to overcome distance limitations (300m over fiber), and extend distance to 40 Km

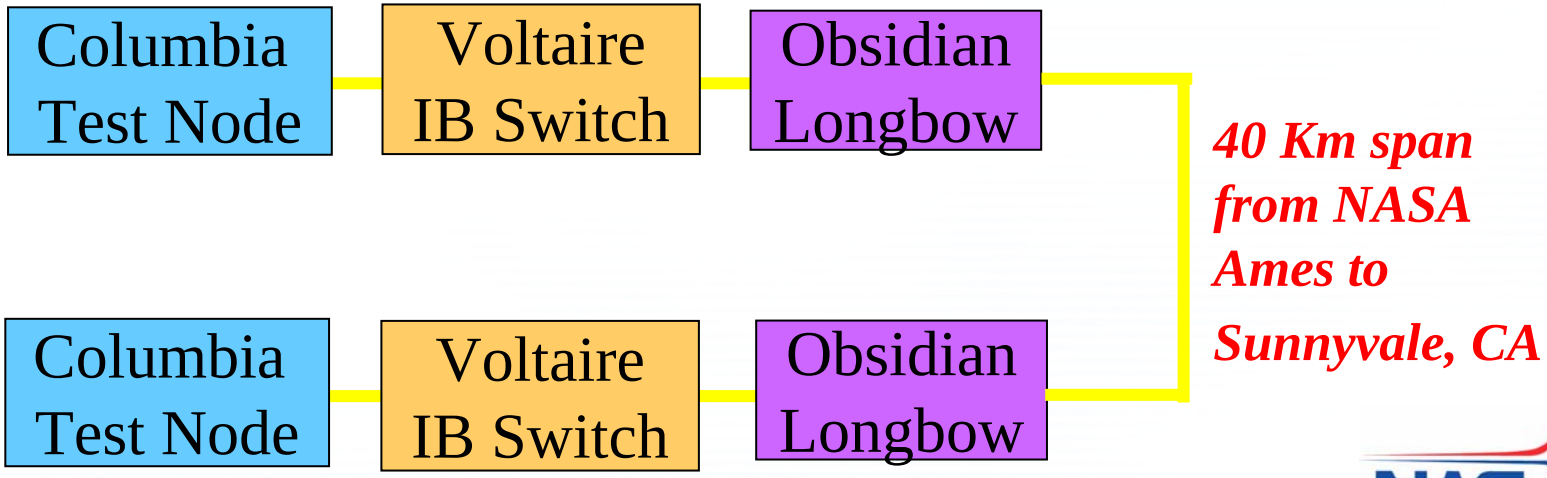

Infiniband Research Results

Results:

- IB in the local environment, produced better performance than traditional Ethernet technology for applications that require extreme low latency

Future:

- Connect Columbia & Hyperwall via longbow to Sunnyvale to demonstrate remote visualization
- Planned demonstration at SC06

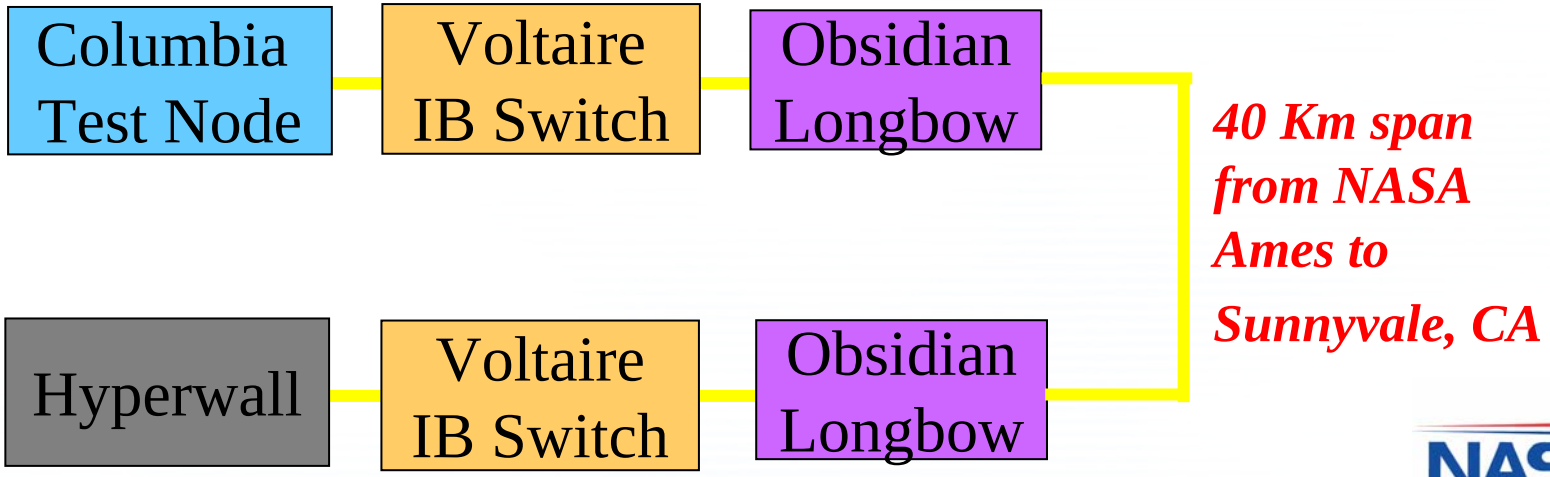

Lambda Agile Switching Research

Goals:

- Demonstrate software-controllable lightwaves, or lambdas, through the use of software configurable optical switching equipment.
- Develop experience in the use of multiple wavelengths per fiber, as one of the mechanisms for making better use of fiber optic resources.

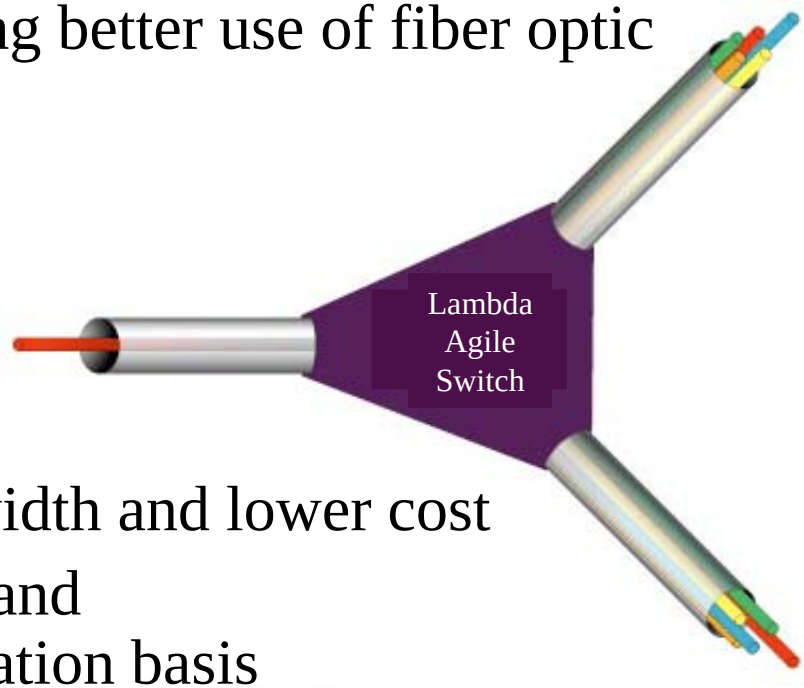

Motivation:

- Using dark fiber resources and customer-controlled optical equipment to increase WAN bandwidth and lower cost
- Ability to enable multiple lambdas and allocate them on a project or application basis

Lambda Agile Switching Research

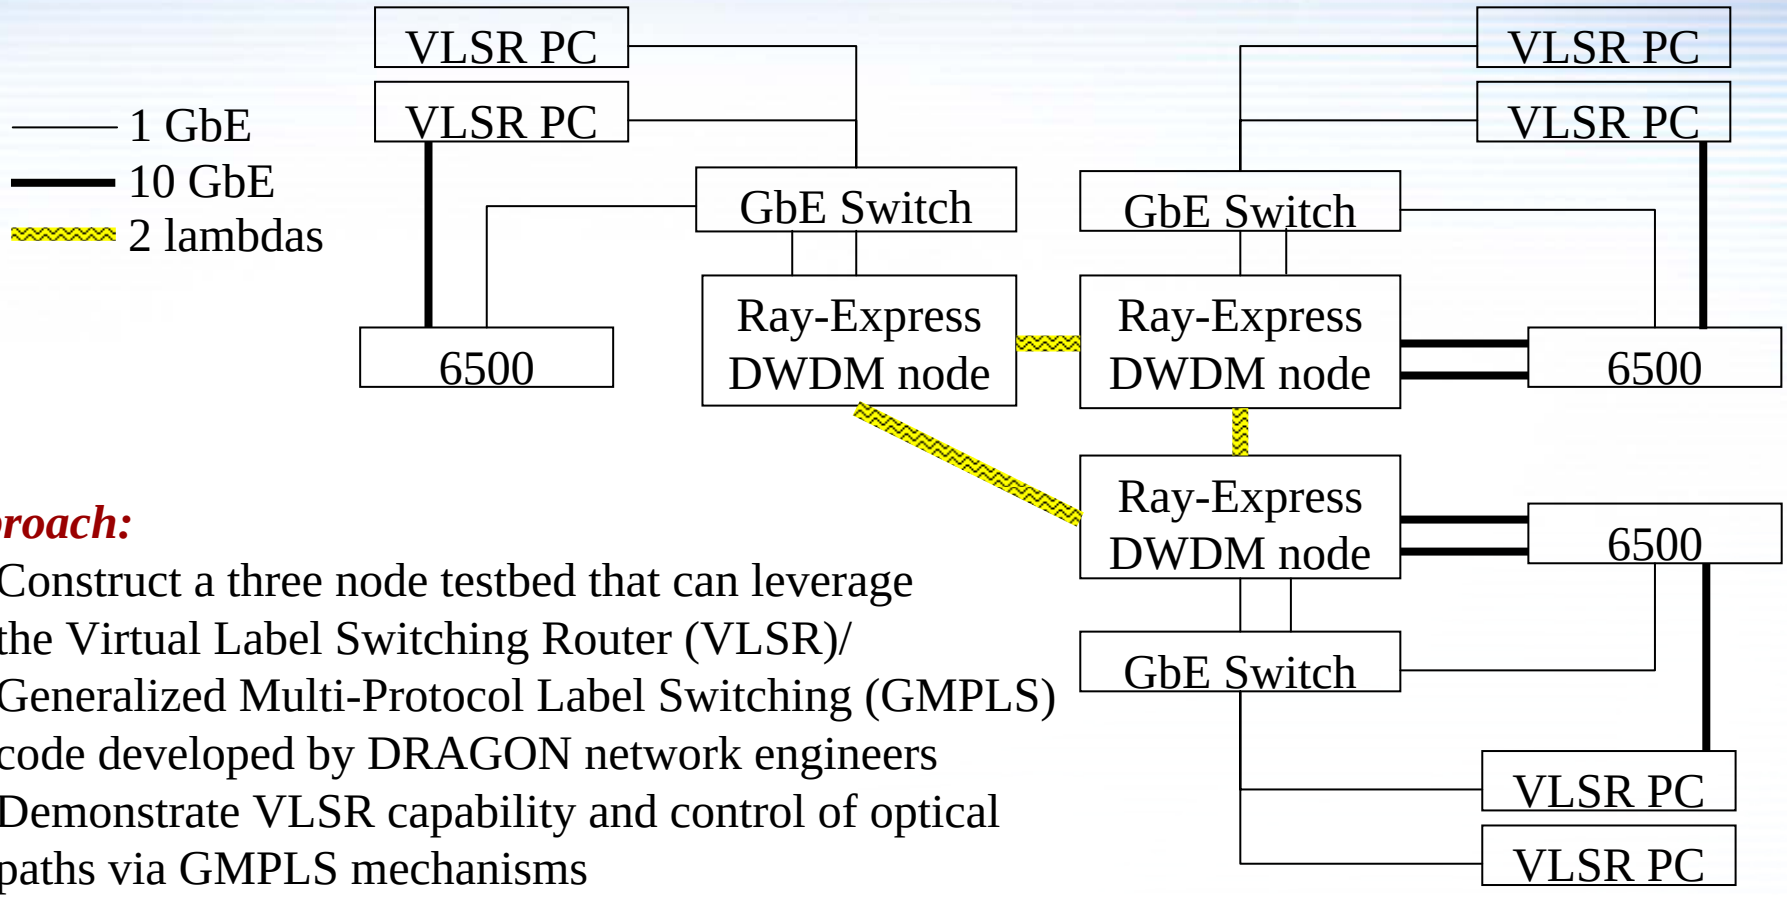

Approach:

- Construct a three node testbed that can leverage the Virtual Label Switching Router (VLSR)/ Generalized Multi-Protocol Label Switching (GMPLS) code developed by DRAGON network engineers
- Demonstrate VLSR capability and control of optical paths via GMPLS mechanisms
- Develop plan to demonstrate dynamic WAN lambda switching
- Collaborate with DRAGON network and other GMPLS-capable networks in support of inter-domain lambda switching

Lambda Agile Switching Results

Status:

- Configured three Movaz RayExpress DWDM shelves to support multiple lambdas with support for 10 Gbps and 1 Gbps interfaces.
- Integrated high performance Linux hosts that can both run the DRAGON VLSR/GMPLS code and perform near line rate throughput at 10 Gbps.
- Established VLSR functionality on testbed switches.

Future:

- Demonstrate GMPLS control of optical paths.
- Provide plan for WAN lambda switching demonstration.
- Invigorate interaction with related projects: HOPI, DRAGON, and DOE ESNet/UltraScienceNet
- Procure multi-degree Reconfigurable Optical Add-Drop Mux (ROADM) for integration with existing testbed equipment, expanding the functionality in support of operational deployment

Enabled Research

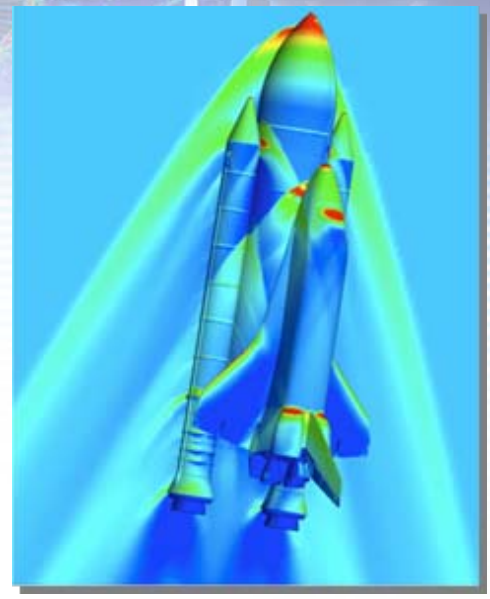

Return to Flight

- CFD techniques are developed and applied to many NASA aerospace analysis & design problems - Full Space Shuttle Launch Vehicle configuration. The results from the computations on Columbia were a critical factor in returning the shuttle to flight.

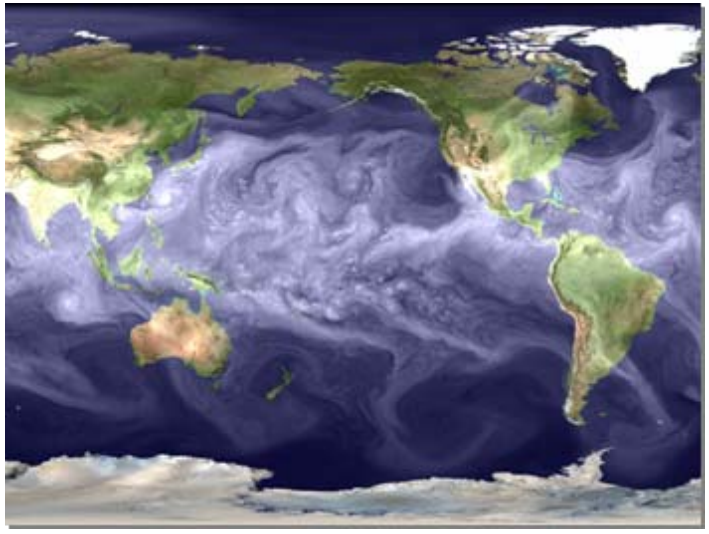

Hurricane Modeling

- NASA researchers used the finite-volume General Circulation Model (fvGCM) to generate real-time, detailed weather predictions and were very successful during the active 2004 and 2005 Atlantic hurricane seasons, providing up to five-day advanced landfall forecasts with a high degree of accuracy for intense hurricanes including

  Katrina and Rita.

Enabled Research

Propulsion Subsystem Analysis:

- High-fidelity unsteady flow simulation techniques for design and analysis enhance safety of propulsion systems for NASA missions and reduces cost of space flight. Computational models of enable design decisions that improve efficiency, performance, and reliability.

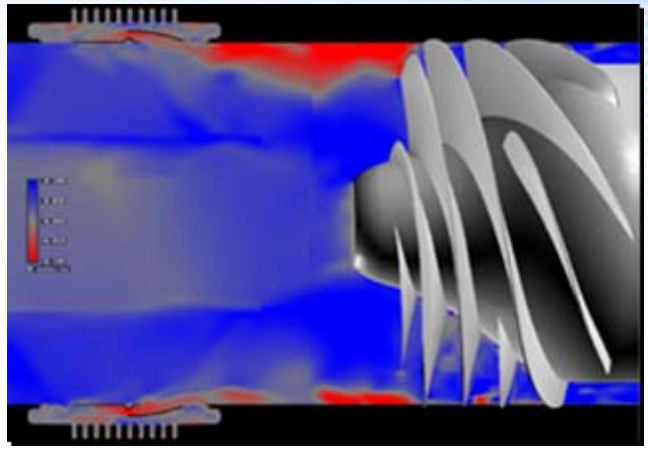

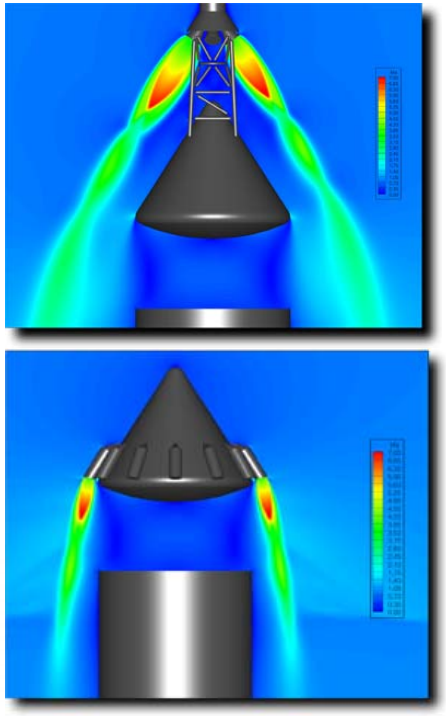

Crew Exploration Vehicle (CEV):

- Supplied high-fidelity simulation data to help evaluate the feasibility of two Launch Abort System (LAS) concepts for the CEV. Compared performance characteristics of these two LAS concepts across different abort scenarios

End-to-End Network Support

- Working directly with the HEC users
  - Documenting environment (transfer application, system, network path) to better assist in trouble shooting performance problems
- Working with local infrastructure staff
  - System Administrators, LAN Engineers, Security/Firewalls, to enable optimal configurations
- Networking testing and analysis to identify and help overcome bottlenecks
- Providing resources for optimal system and application configurations for efficient use of data transfer applications: scp, bbftp, etc.

Additional Information

- NASA Ames Research Center
  http://www.arc.nasa.gov
- NASA Advanced Supercomputing Division (NAS)
  http://www.nas.nasa.gov
- NASA Research and Engineering Network (NREN)
  http://www.nren.nasa.gov
- National Lambda Rail (NLR)
  http://www.nlr.net
- Pacific Wave
  http://www.pacificwave.net
- DRAGON
  http://dragon.east.isi.edu

Thank You!

ONT3 Presentations
will be available
on the web
by
Monday, September 18th

http://www.nren.nasa.gov/workshop9/