UltraScience Net

Ultra High-Speed Research Network for Large-Scale Science

Bill Wing, Nagi Rao

Oak Ridge National Laboratory

Historical Perspective and Motivation

- Early adoption of advanced network technologies: Examples: ESnet was first to WAN deployment of ATM; TCP slow start and congestion control
- NGI QOS Testbed – First network-wide testbed for QoS technology used in the Qbone
- Clipper Project – High-speed data transfer testing of dedicated OC-12 for High Energy applications
- MORPHnet: A framework for using a fraction of ESnet bandwidth to conduct experimental networking activities - never implemented
- Experimental (breakable) networking environment to develop, test, deploy, and prototype advanced networking technologies and network-intensive high-impact applications is a critical component in SC networking roadmap
- Advanced cost-effective networking technologies for science communities beyond the Internet are not commercially available
- The unique networking requirements of the DOE Office of Science are not adequately and timely addressed by the industry and other federal network research programs
- Budget trends and realities require innovations and new models of cost-effective networking technologies

Series of DOE Workshops

Development of Office of Science Networking Requirements

DOE Networking Workshops:

- **High-Performance Networks for High-Impact Science, Aug 13-15, 2002.**
- **Network Provisioning and Protocols for High-Impact Science, April 10-11, 2003.**
- **DOE Science Networking Challenge: Roadmap to 2008, June 3-5, 2003**

Other Agency Workshops:

1. NSF Workshop on Ultra-High Capacity Optical Communications and Networking, October 21-22, 2002
2. NSF Workshop on Network Research Testbeds, October 17-18, 2002

DOE Networking Workshop Summary - I

- Diverse domain networking requirements
  1. Guaranteed QoS, best-effort, real-time capabilities, batch services, email, local/national/intercontinental large file transfers
  2. On-demand bandwidth, dedicated bandwidth, shard network capabilities, SANs-LANs challenges
  3. Transparent cyber-security, distributed trust, Lab-based cyber security architecture, DOE cyber security policy, etc.

- Distributed terascale computing facilities and petabytes data archives need seamless access
  1. Resources located across the country
  2. Users located at universities, national laboratories, industry

- Critical and timely networking requirements for large-impact science
  1. Advanced and deployable networking capabilities for LHC experiments, SNS, climate modeling, astro-physics, and computational biology

  2. Leverage current opportunities in telecommunication industry and mature optical network technologies to build advanced networking infrastructures for science

DOE Networking Workshop Summary - II

Science Areas	Today End2End Throughput	2008 End2End Throughput	2013 End2End Throughput	Remarks: Basic research, testing and deployment
High Energy Physics	0.5 Gbps E2E	100 Gbps E2e	1.0 Tbps	high throughput
Climate Data & Computations	0.5 Gbps E2E	160-200 Gbps	*n* Tbps	high throughput
SNS NanoScience	does not exist	1.0 Gbps steady state	Tbps & control channels	remote control & high throughput
Fusion Energy	500MB/min (Burst)	500MB/20sec (burst)	*n* Tbps	time critical transport
Astrophysics	1TB/week	N*N multicast	1TB+ & stable streams	computational steering & collaborations
Genomics Data & Computations	1TB/day	100s users	Tbps & control channels	high throughput & steering

Office of Science Networking Roadmap
(2003 Workshop Report)

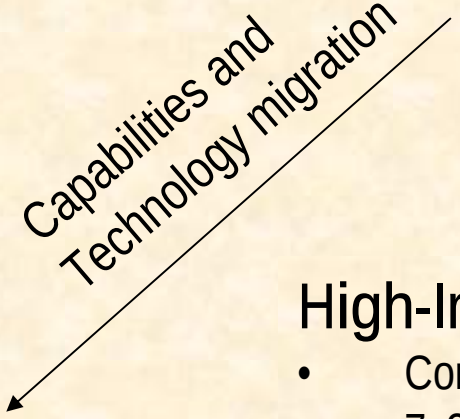

Advanced Research Network
- R&D – Breakable components
- Scheduled operations
- Ultra high speed components
- Electro/optical components

High-Impact Science Network
- Connect few science sites
- 7x24 operations
- Very high speed – four 9s
- Specialized components

Production Networks
- Connects all DOE sites
- 7x24 & reliable – four 9s
- Advanced Internet capability
- Predominantly best-effort

DOE Network Research Program

for Large-Scale Science

Vision

- Develop, test, and deploy advanced network technologies to efficiently and securely interconnect scientific resources such as distributed terascale computing resources, petabytes-scale data archives, remote visualizations centers, complex scientific instruments, and research teams

Program Elements

- Research, Development, and Engineering
  - **Foundation of high-capacity networks**
- Advanced Network Research Testbeds (ANRT) (24% of FY04 budget)
  - ✓ **Deployment, and testing of advanced network technologies -- UltraScience Net**
- Workforce development
  - Young investigation program

Networking for Large-Scale Science:
UltraScience Net and ESnet span only the core

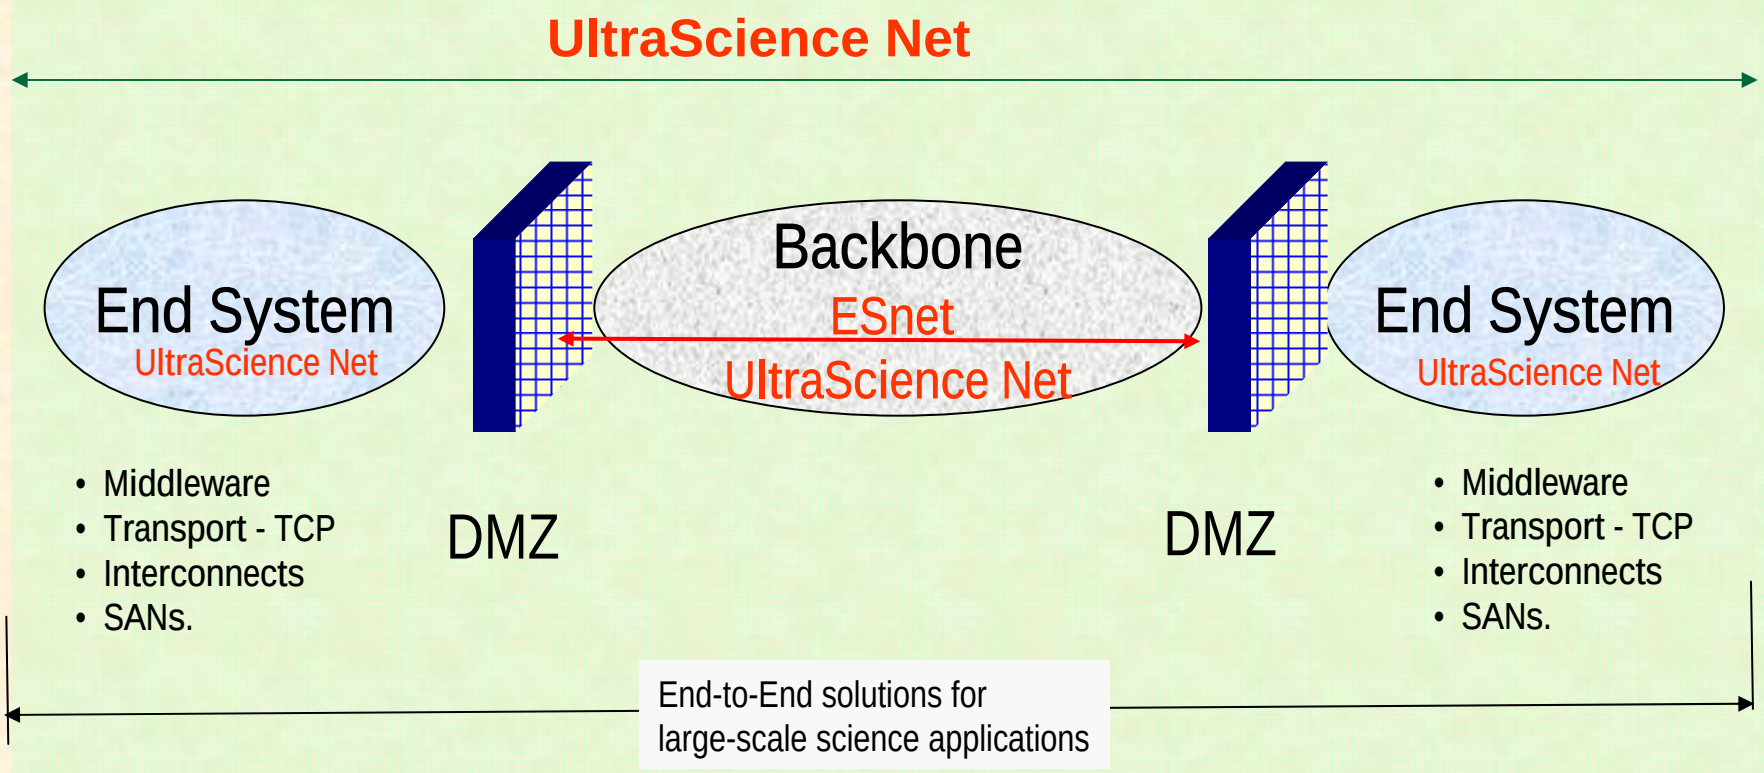

Current Areas of Network Research

- Ultra high-speed data transfer protocols – TCP enhancements and alternatives for ultra high-speed networks
- Dynamic provisioning of DWDM optical networks - on-demand bandwidth and dedicated channels
- Network measurement and analysis – end-to-end performance monitoring, prediction, and fault diagnosis
- ✓ Advanced research networks – experimental network prototyping, testing, and deployment
- Scalable cyber security systems – Firewalls and intrusion detection systems

The Case for UltraScience Net

- Meeting the diverse network requirements for high-performance and agile networking beyond current best-effort networks (Internet and ESnet)
- Lack of investment by industry on cost-effective advanced networking technologies for large-scale science
- Implementation of Office of Science recommendations in network roadmap to 2008
- Flat networking budget (FY03, FY04, FY05, FY06?) → Innovation and optimization of existing resources
- Enables development and testing of advanced network capabilities for critical science projects such as LHC (ATLAS and CMS), computational genomics, and SNS, etc.
- Opportunity in industry: abundant dark fiber and mature DWDM optical technologies

UltraScience Net

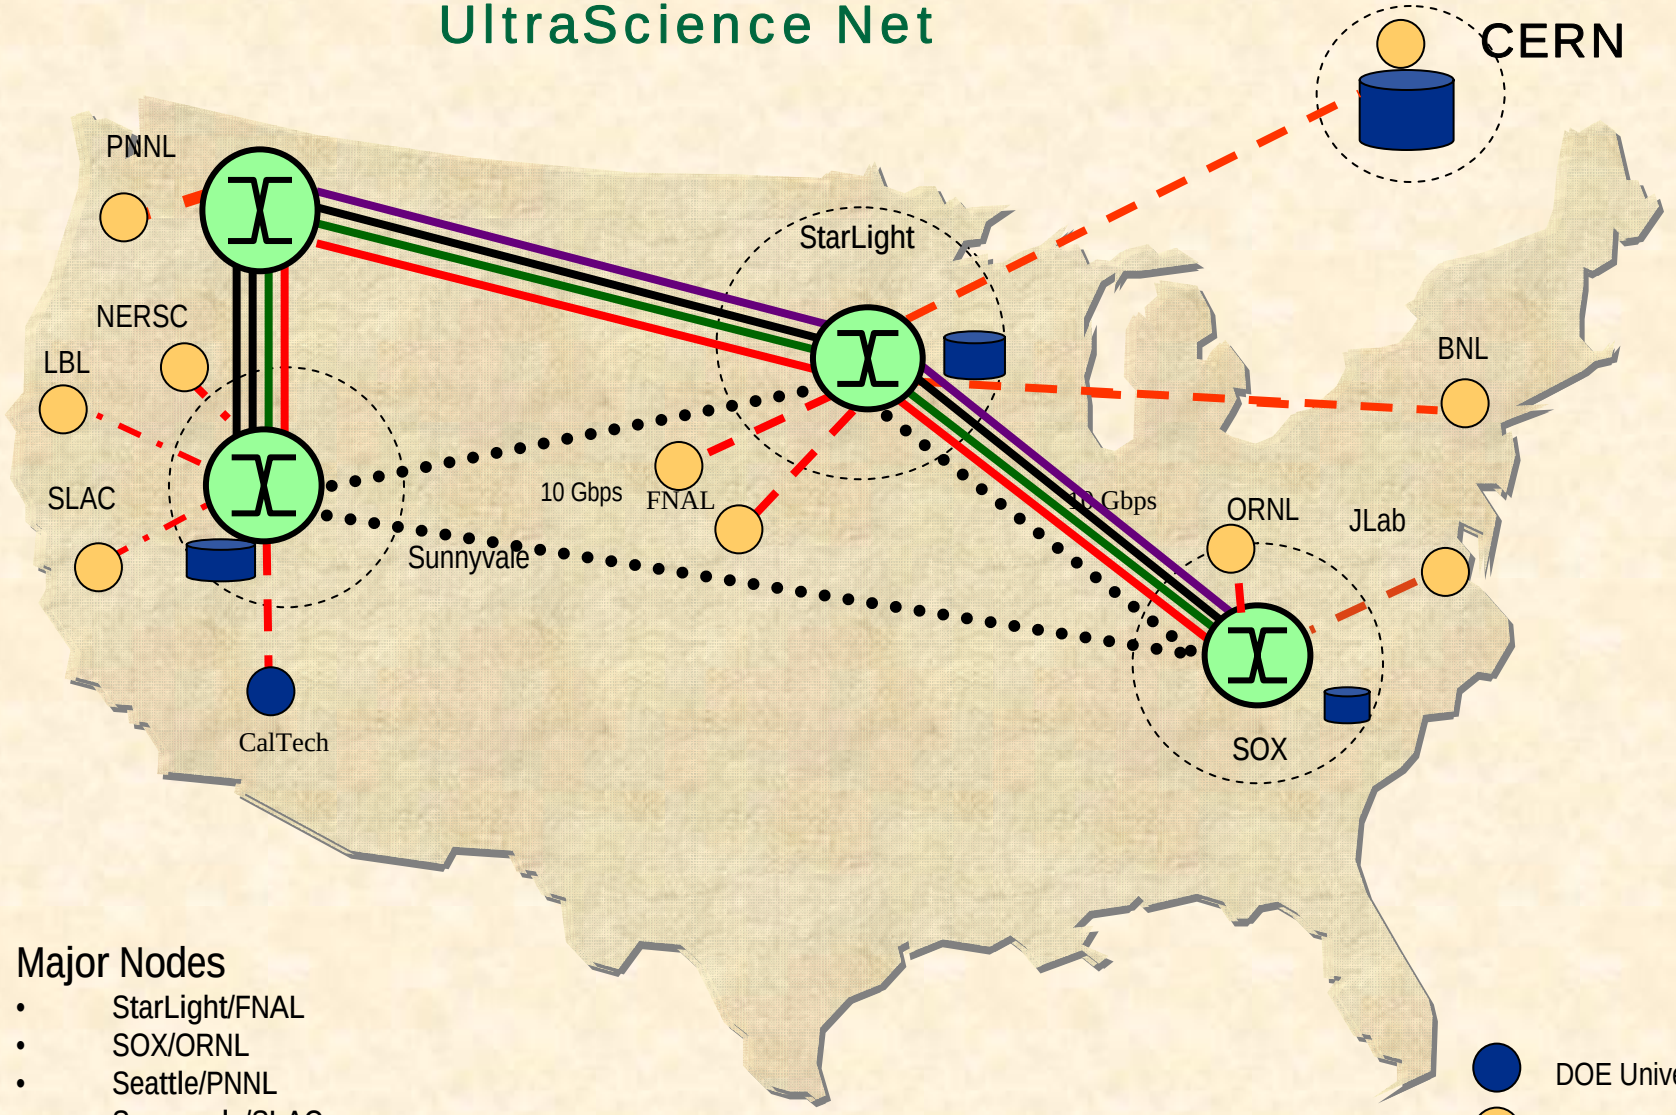

Major Nodes

- StarLight/FNAL
- SOX/ORNL
- Seattle/PNNL
- Sunnyvale/SLAC
- Sunnyvale/Caltech

10 Gbps ESnet Links

10 Gbps UltraScience Net Link

under discussion

DOE University Partners

DOE National Lab

OAK RIDGE NATIONAL LABORATORY
U. S. DEPARTMENT OF ENERGY

UT-BATTELLE

Technologies to be developed on UltraScience Net

- On-demand dedicated bandwidth channels
  1. Dedicated SONET/Ethernet channels to science applications – IP/non-IP
  2. Dedicated wavelengths to science application - IP/non-IP
- Control/Signaling plane for dynamic provisioning
  - TL1/GMPLS
  - Secure end-to-end signaling
- Ultra high-speed transport protocols testing
  1. High-throughput TCP variants
  2. TCP alternatives for IP/non-IP
- Prototyping science application for ultra high-speed networks
  1. High-speed data transfers – High energy physics (CSM and ATLAS) ** SciDAC
  2. Remote computational steering – Nuclear and astro physics (TSI) **SciDAC
  3. Remote visualization – Computational genomics
  4. Remote Instrument control - Biology
- Ultra high-speed network components
  1. Cyber security – packet filters and firewalls
  2. Network measurement – network and host levels

Community Support for UltraScience Net

DOE Labs

- FNAL — Provide fiber connection to UltraScience Net at Starlight
- ORNL — Fiber connection to UltraScience Net in SOX-Atlanta
- SLAC — DWDM connection to Sunnyvale-California (under discussion)
- PNNL — Fiber connection to Seattle, Washington

Proposals from science communities to use UltraScience Net

- High Energy — FNAL - High-throughput transfers of LHC data
- Genomics — PNNL - Remote visualization of computational biology
- Astrophysics — ORNL - Real-time data visualization and steering
- Fusion — General Atomics –WAN QoS using MPLS for fusion energy application
- LBL — Exploration of MPLS technology for ESnet

Inter-Agency - Collaboration

- CHEETAH — NSF testbeds use UltraScience Net to reach east coast
- DRAGON — NSF all-optical network to UltraScience Net to reach Starlight in Chicago
- HOPI — Internet2 traffic exchange with UltraScience Net

Technology Transfer: UltraScience Net to ESnet

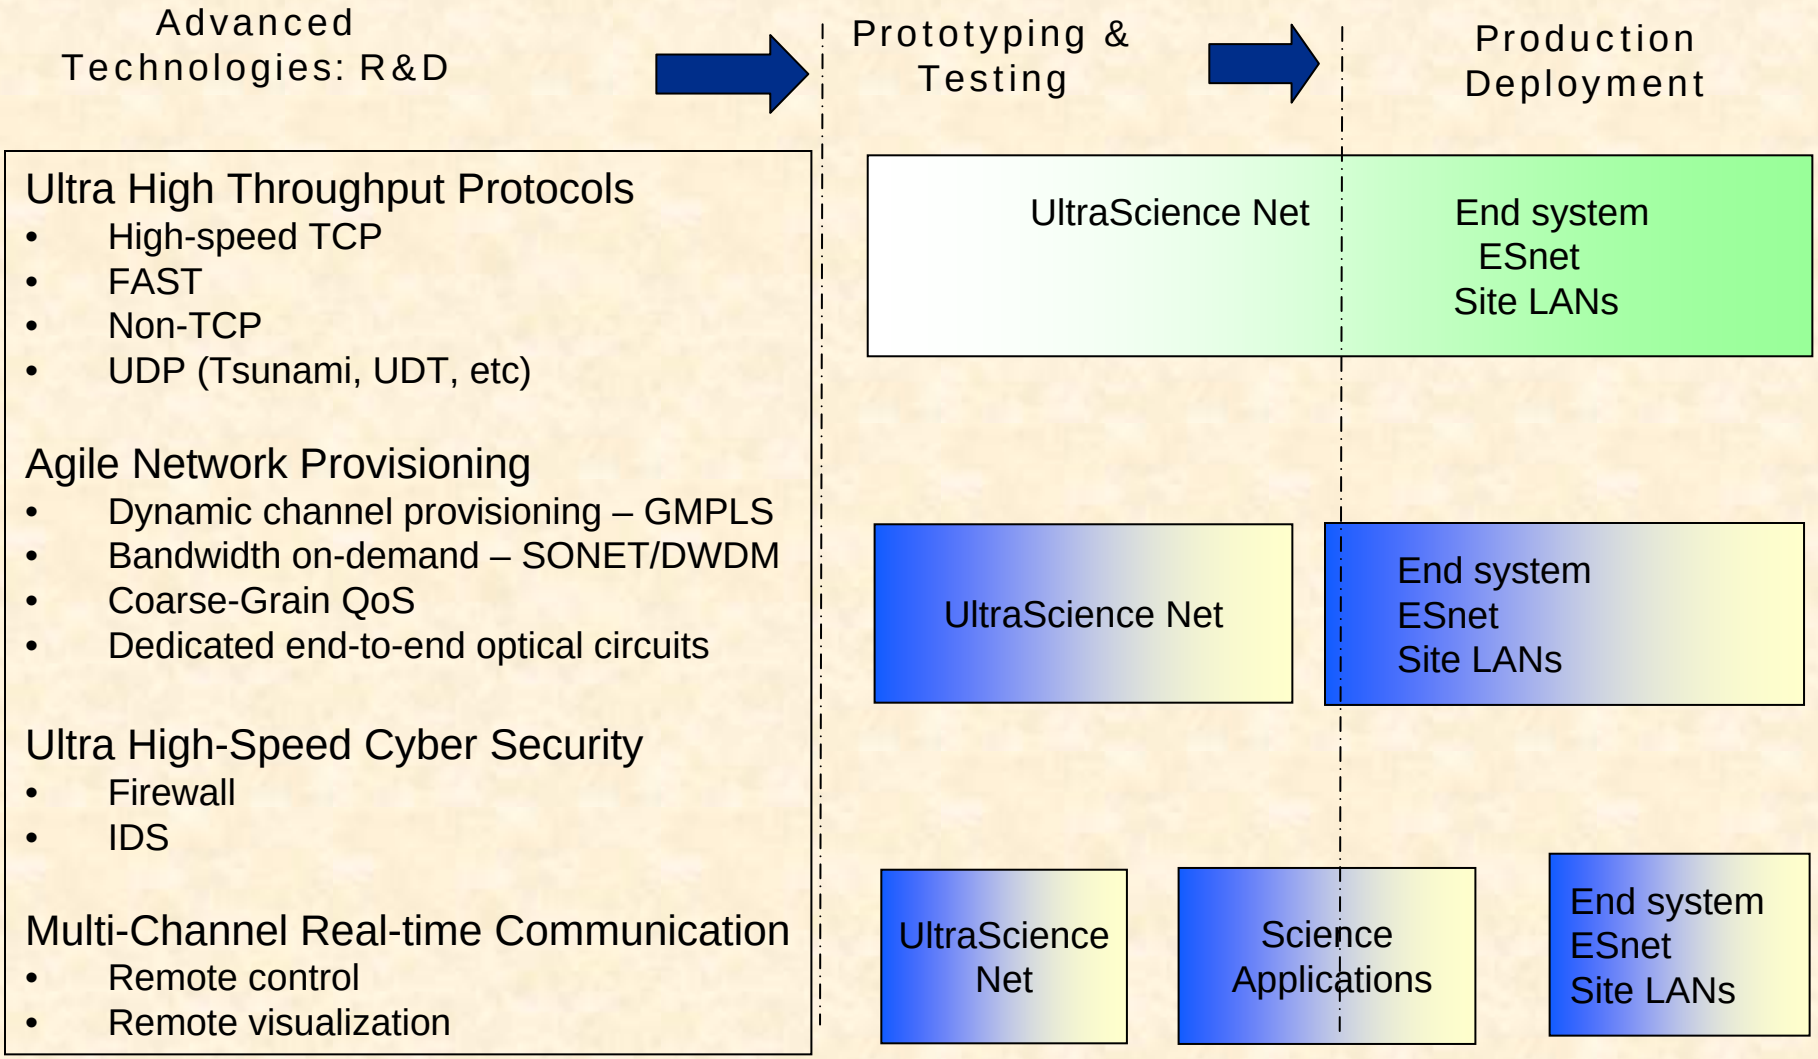

UltraScience Net and ESnet: Combined efforts to meet large-scale science challenges

UltraScience Net Features

- Experimental, breakable, unreliable
- Connect selected DOE sites and research networks
- Scheduled experiments on network research and application prototyping
- Dynamic provisioning of dedicated channels
- On-demand QoS
- Application to-application

UltraScience Net Technologies

- **Switching/Provisioning**
  Circuit switching and hybrid packet/circuit provisioning
- **Backbone**
  On-demand provisioned SONET/DWDM, 10GigE LAN PHY
- **Control Plane**
  Out-of-band TL1 and in-band GMPLS
- **Transport**
  1. TCP, UDP and other
  2. non-TCP for dedicated channels

ESnet Features

- Production, highly reliable – 4 9s
- Connect all DOE sites and universities
- 7x24 availability for all services provided to all sites
- Static provisioning
- No QoS: Best-effort
- Edge-to-edge

ESnet Technologies

- **Routing**
  Packet/routed network
- **Backbone**
  POS – packet over SONET
- **Control Plane**
  Not applicable
- **Transport**
  1. TCP, UDP and others

ESnet and UltraScience Net Activities

Engineering Collaboration

- Joint ESnet/UltraScience Net engineering design team for UltraScience Net
- ESnet is a collaborator of UltraScience Net
- ESnet engineering participate in the panel review of ultra network/application proposals
- Annual joint ESCC and network research PIs meetings
- ESnet MPLS roll-out on Sunnyvale-Chicago in support of scheduled bandwidth

Inter-Agency Activities Involving ESnet and UltraScience Net

ESnet:

LSN/JET member
LSN/JET co-chair (G. Seweryniak)

UltraScience Net

LSN/NRT Member
LSN/NRT co-chair (T. Ndousse)
NSF CHEETAH (Rao, Wing PIs)

UltraScience Net Operations and Management

Engineering Team

1. UltraScience Net Engineering
2. ESnet Engineering
3. Application Developers

Management Team

- UltraScience Net Engineering
- ESnet Engineering Rep
- ESCC Rep
- ESSC Rep

Research Team – Awards Pending

1. Network Research PIs
2. Application Prototyping PIs

Management Team Responsibilities

1. Prioritize experiments on UltraScience Net
2. Schedule testing
3. Develop technology transfer strategies

Summary

- UltraScience Net is an advanced experimental network testbed
  - demanded and supported by large-scale science applications of Office of Science
- Large-scale science applications require extreme networking
  - requires in-house efforts because it not cost-effective for industry
- UltraScience Net is an integrated infrastructure
  - developing, testing, and deploying advanced network technologies for next-generation science applications

Q&A

UltraScience Net Website: http://www.csm.ornl.gov/ultranet