Final Report & Highlights from the Subcommittee for the Examination of the Role and Efficiency of Networking and Networking Research within the Office of Science

Ellen B. Stechel, Presenting

26-February-2008

The Charge:

"weigh & review the organization, performance, expansion, and <u>effectiveness</u> of the current operations of ESnet. .. consider the <u>proposed evolution</u> of ESnet, its appropriateness and comprehensiveness in addressing the <u>data communication needs</u> .. that will <u>enable scientists nationwide to extend the frontiers of science</u>. .. make suggestions and recommendations on the <u>appropriateness and comprehensiveness of the networking research</u> .. with a view towards meeting the <u>long-term networking needs</u> .."

The Distinguished Sub-Panel

Last Name	First Name	Affiliation
Bailey	Ron	ASCAC Member
Corones	Jim	Krell Institute
Foster	Ian	ANL & U. of Chicago
Hitchcock	Dan	OASCR Liaison
Huntoon	Wendy	Pittsburgh Supercomputing Center (PSC)
Newman	Harvey	California Institute of Technology (Caltech)
Rahn	Larry	SNL & DOE-BES-Chemical Sciences
Simon	Horst	ASCAC Member
Sobieski	Jerry	Mid Atlantic Crossroads (MAX)
Stechel	Ellen	Chair & ASCAC Member
Wing	Bill	Co-Chair & ORNL

Interpreting the Charge

- Cornerstone and Framing:
  - **"that will enable scientists nationwide to extend the frontiers of science."**
- Other Key Phrases
  - With respect to ESnet
    - ***Expansion and effectiveness of current operations, Proposed evolution***
    - In light of the Lehman review and more recent Operational Review, this is de-emphasized and the subpanel's report concurs with those findings and recommendations.
  - With respect to needs
    - ***Data communication needs – pulled this into the title***
    - ***Long-term Networking needs***
      - **Broad interpretation in considering how the network will affect all major areas of future cyber-infrastructure and its anticipated role in enabling new and transformative science processes over the coming decade.**
  - With respect to Networking Research to meet long-term needs
    - ***Appropriateness and comprehensiveness of the Networking Research***
    - Focused on motivating why the office should re-invigorate its networking research program
- Timeframe – **5-10 Years**
  - The network technologies and services considered will require substantial and sustained efforts to realize
- Cyber Security –
  - Expect this will be very important, important enough that it will be and is already the focus of independent efforts
  - To avoid duplication, this report does not have findings nor recommendations with respect to Cyber Security

Report Title

Data Communications Needs:
Advancing the Frontiers of Science
Through Advanced Networks and Networking Research

Timeline

- Panel Reformulated – Completed Jan 2006
- Teleconferencing, Web Meeting Capability, and Website Established – January 2007
- Begin Teleconferences in February
- Status Report – Delivered to ASCAC Tues 27-Feb-2007
- ~Bi-Weekly Meetings continued
- ESnet Visit (13-April-2007)
- Status Report – ASCAC 14-Aug-2007
- ~Bi-Weekly Teleconferences Continued
- DRAFT Preliminary Report – Distributed to ASCAC – 02-Nov-2007
- Preliminary Report – Delivered to ASCAC 6-Nov-2007
- Writing Session in Chicago – 13-Dec-2007
- Final Report – Distributed to ASCAC 15-Feb-2008
- Final Report Discussion – ASCAC 26-Feb-2008

Six Baseline Observations/Trends

First: The network now underlies nearly every aspect of advanced e-science and, in fact, nearly all science activity.

Second: Major Drivers for Advanced "Cyber" Resources

- There is general acceptance that computational ***simulation and modeling*** has become the "third leg" of scientific inquiry. As such, it constitutes a major driver for cyber resources.
- Beyond the basic computational requirements, these applications depend upon tightly integrated access to storage facilities, visualization facilities, and advanced high-performance networks.

Third: Accelerating Trend driving the Need for Data Packaging, Access, and Mobility

- Scientific inquiry is moving towards more quantitative understanding of ever more complex systems,
  - whether that is from sophisticated experiments in SC's one-of-a-kind facilities
  - or from large-scale modeling and simulation.
- Acceptance or refinement of the conclusions from such scientific inquiries depends on scrutiny from many eyes and from different research groups, approaches, and perspectives.
- This implies a need for data packaging, access, and mobility.

Fourth: Fundamental Changes in Data Distribution and Use

- There is an accelerating trend in the quantity and diversity of raw data
  - captured from instruments, computations, and sensors and
  - archived for reuse,
- again implying a strong driver for data handling, accessibility, and mobility.

Fifth and Sixth: Virtual Facilities, Collaboration, and System of Systems

- The continuing trend is toward global collaborative e-science,
  - where science teams, science facilities, and science itself crosses a wide range of boundaries,
  - including traditionally distinct disciplines, funding agencies, and nations, once again
- indicating a major driver for data communication and data mobility
- An increasing trend toward the development of new cyber systems by combining and integrating existing facilities is driving a need for a System of Systems approach to scientific infrastructure.

System of Systems Typically Exhibit the Behaviors of Complex Adaptive Systems

- Operational Independence of Elements
- Managerial Independence of Elements
- Evolutionary Development
- Geographical Distribution
- Inter-disciplinary
- Heterogeneity of Systems
- System of Networks
- Emergent Behavior

System of Systems Science is an emerging discipline not yet claimed by anyone, but can provide a critical organizing principle

See for example: Carlock, P.G., and R.E. Fenton. "System-of-Systems (SoS) Enterprise Systems for Information-Intensive Organizations," Systems Engineering, Vol. 4, No. 4 (2001), pp. 242-261.

The Office of Science facilities infrastructure represent a System of Systems

- **Complex Systems with Many Similar and Different Components**
  - Producing data at a phenomenal rate
  - Producing large amounts of heterogeneous, geographically dispersed data
  - Relationship and coupling of data and computation
  - Cultural barriers
  - Growing in its dependence on the emergence of advanced networks
- **Systems, Projects, People**
  - Computing System
  - Experimental Facilities, Instruments, and User Centers
  - Observational Networks
  - Data, Information Resources, and Visualization
  - Virtual Organizations and Distributed Collaborations
  - OASCR Research Projects:
    - Applied Math, Computer Science, Networking Research, SciDAC
  - Researchers, Facility Staff, Center Staff, Program Managers
- **The Emergent Behavior is Breakthrough Advances in Scientific and Engineering Research**
  - Accelerating and amplifying impact on urgent problems of global scale

Seven “High Level” Findings

First: About ESnet

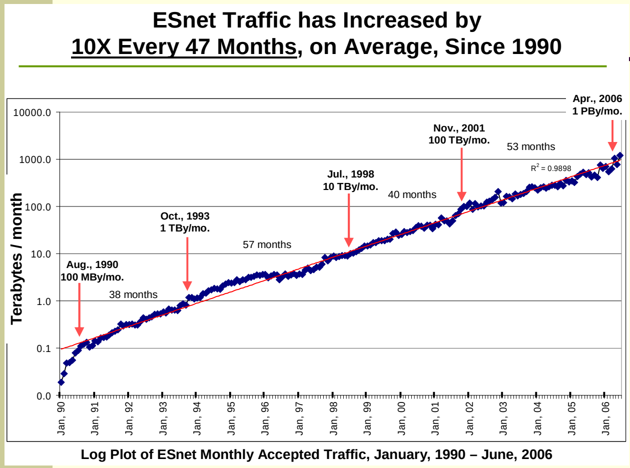

- ESnet facility is doing an exemplary job in architecting, deploying, and operating a high performance network infrastructure to serve DOE science needs.

- The successful basis for the operations has been a continuous “requirements” driven process.
- The subpanel concurs that ESnet4 infrastructure is critical in the relative near term and will play a continuing vital role in the long term as new network architectures and services evolve and mature.
- Sustained and adequate funding for these facilities remains essential.

Second: DOE Science Cannot Depend Solely on the Commercial Sector

- High-performance networking is critical to programs in the Office of Science (SC,) including high-performance computing.
- The high-performance networks required will not automatically emerge from commercial R&D.
- Therefore, SC will need to fund strategic, high-performance networks and networking research.
- A corollary is that high-performance networks and networking research need the same level of attention as given to high-performance computing.
- The challenge then, given the central importance of high-performance networking to SC programs, is to construct a long-term and far-reaching network R&D program that sustains innovation from basic research to prototype to early deployment and culminates in a production network environment for use in advancing science.

Next Three:

- The pursuit of DOE science will continue to advance with an accelerating dependence on networks and related technologies to unite disparate teams, to allow efficient exchange of information and applications, and to enable new modes of scientific inquiry.
  - The challenge is to develop network architectures and service models that support virtual organizations and provide reliable, predictable, and repeatable network performance, accessibility, and security.
- A key movement that will have profound impact on science, society, and the economy in the coming decade will be the development of national and global-scale, data-intensive (terascale to petascale and beyond) distributed "cyber environments."
  - The challenge will be to develop dynamic and intelligent (cyber) resource allocation architectures that will allow the DOE scientist to transparently and easily make use of resources at this scale.
- The network capacity and service capabilities anticipated for the next decade of science activities will likely be three to four orders of magnitude greater than current network architectures and technologies can effectively address today.
  - The challenge will be to accelerate or develop a ten-year "technology" trajectory to achieve this projected need.

Final Two: The Valley of Death and Domain Specific Advances

- There are many barriers to the development and adoption of promising new networking research ideas (across all agencies and industry.)
  - These range from too narrow a focus (e.g., driven by already evident deficiencies,) to the classic “valley of death” between pure research and robust documented product.
  - In the DOE environment, long-term delays in finding and leveraging useful new network technologies hinder the emergence of new scientific systems.
  - The challenge is to establish mechanisms to identify promising networking research concepts and move them progressively through prototyping, experimental deployments, and ultimately into a production network service environment.

- Out of necessity, a few leading-edge science communities have pushed the state of the art in networks and networking services.
  - The result has been vast discrepancies in the level of capabilities for data distribution and management available across the scientific enterprise.
  - The challenge is twofold:
    - first, to leverage and generalize the services developed in one context for the broader utility of the DOE science community and
    - second, to develop advanced cyberinfrastructure service architectures that allow future efforts to create, easily and effectively, the types of network (or other cyberinfrastructure) that they need.

A Strategy for Moving up the Maturity Continuum and Spanning the “Valley of Death”

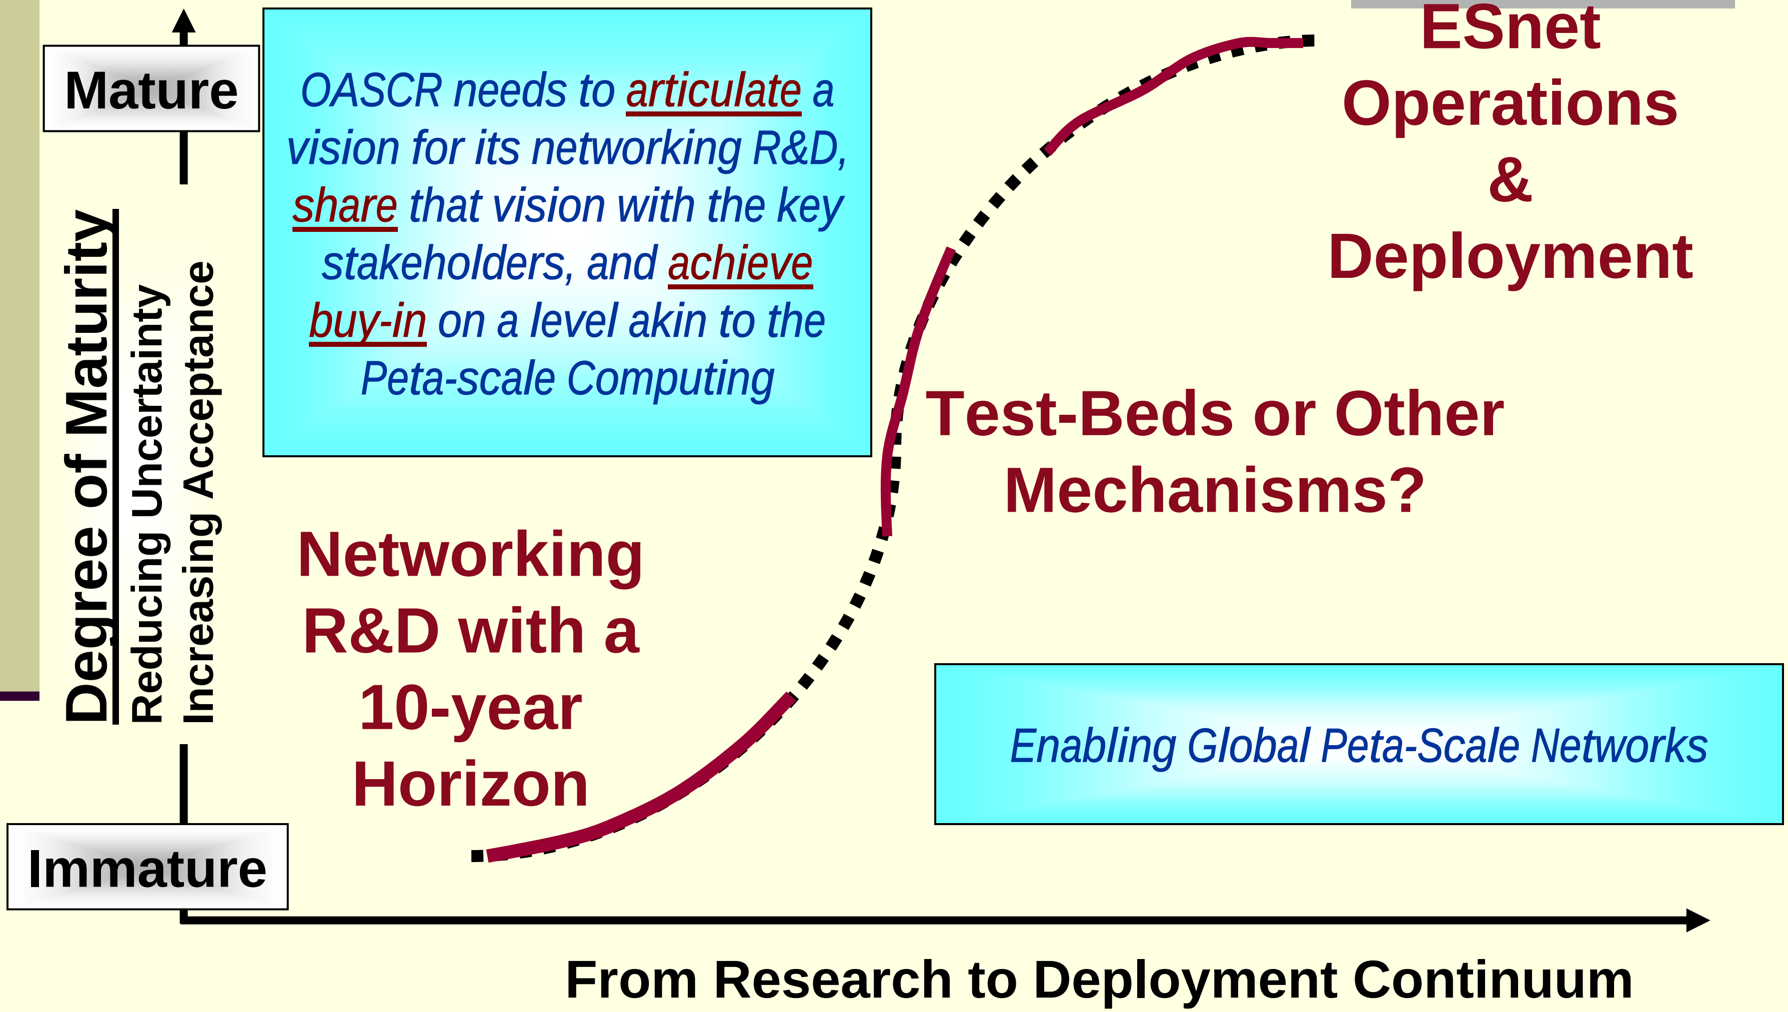

Six “High Level” Recommendations

Recommendation #1

- Approach the development of advanced networking in a fashion similar to the manner in which SC develops its goals and objectives for petascale and exascale computing.
- The **network and networking services need to be an explicit and fundamental element of advanced petascale and exascale science,** of equal importance to leadership-class computing and large-scale scientific facilities in **enabling high-end science.**
- Hence, the subpanel encourages SC to create mechanisms whereby the DOE science community is encouraged to think broadly about how an unconstrained network resource might transform the conduct of science.

Recommendation #2

- Establish mechanisms to manage the implications and issues that arise from the **System of Systems** aspects of DOE science, facilities, and programs.
- Such an approach will increasingly be required to enable interoperability needed for multi-disciplinary science.
- Advanced information technologies will enable this System of Systems; advanced networking architectures and capabilities will be particularly critical.

Recommendation #3

- Reinvigorate an **aggressive, sustainable, long-term, and strategically focused networking research and development program** to create network-specific technologies that will allow DOE not just to increase the speed of existing systems but also to **transform the manner in which science is done**.
- The recommendation is that ASCR should convene an external committee to review their networking research program on a regular basis, in order to maintain a ten-year research horizon and integration across SC.

Recommendation #4

- Formulate a deliberate **strategy to bridge the "valley of death"** for a networking research program that helps advance DOE science.
- Such a strategy would move concepts (whether from within ASCR's basic networking research program or from the networking research community more generally) through testbed deployment to production.
- The "valley of death" issue could be addressed with funding to support applied R&D and experimental or early adopter deployments.
- The strategy should involve increasing collaboration between the networking research scientists and operational/engineering facilities' personnel as concepts mature.
- Moreover, the application science community, as end-users, must be a continual and integral component of the collaboration.

Recommendation #5

- Explore and develop methods for the automated monitoring, troubleshooting, diagnosis, and management of advanced network architectures and the applications that operate on those networks.
- The effective use of **next generation networks**, with unprecedented and rapidly expanding capacity and utility as well as complexity, **will require a new paradigm of operations and management**, including:
  - end-to-end monitoring of the network (necessarily including the end systems;)
  - automated methods for operating, managing, diagnosing, and alerting;
  - and use of higher-level services to ensure optimal network resource utilization and workflow coordination and management.
- ASCR could convene a workshop to flesh out these and other priority research directions.

Recommendation #6

- Integrate research in **data collection, archiving, curation, generation, pedigree, and access** into the ASCR networking research program.
- Effective integration of new networking technologies into the future data-management architecture is crucial to providing timely and secure access to and efficient migration or access to large datasets across an emerging cross-disciplinary globally, distributed science environment.

Questions/Discussion

Thank-you for your Attention