Cybersecurity for HPC Systems:
State of the Art and Looking to the Future

Sean Peisert
Lawrence Berkeley National Laboratory

NIST HPC Security Workshop — March 28, 2018

U.S. DEPARTMENT OF ENERGY

UNIVERSITY OF CALIFORNIA

BERKELEY LAB
Lawrence Berkeley National Laboratory

Why are we here?

HPC centers are targets — what does that mean?

What Are the Threats to HPC?

- Confidentiality
  - Data leakage (even in “open science”)
- Integrity
  - Alteration of code or data
  - Misuse of computing cycles
- Availability
  - Disruption/denial of service against HPC systems or networks that connect them

Who are the Attackers?

- Sometimes ***external attackers***…

- Sometimes ***insiders***

  - ***Insider*** — someone who has some combination of:
    - ***access*** to a resource,
    - ***knowledge*** of an organization, and/or
    - ***trust*** by an organization.

    - ***There can be degrees of this.***
      - System administrators
      - External, authorized HPC users

These Threats Are Not Just Theoretical…

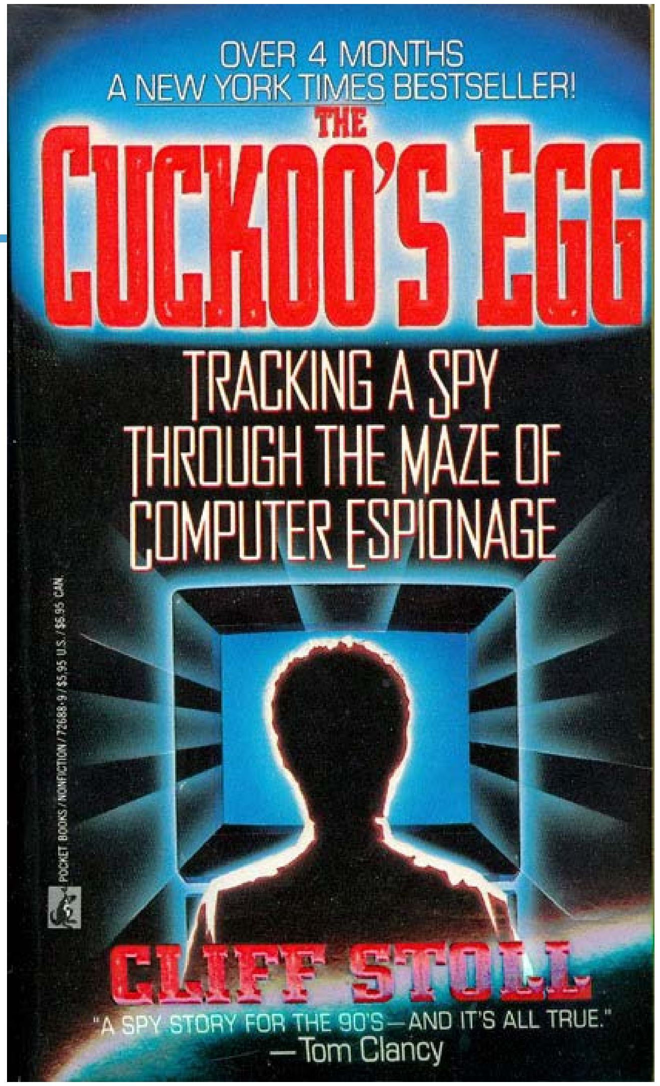

- "Wily hacker" who broke into DOE and DOD computing systems in the mid-1980s.
  - C. Stoll, "Stalking the Wily Hacker," *Communications of the ACM*, 31(5), May 1988.
- "Stakkato" attacks against NCAR, DOE, and NSF-funded supercomputing centers in the mid-2000s.
  - L. Nixon, "The Stakkato Intrusions: What Happened and What Have We Learned?" *Proc. 6th IEEE International Symposium on Cluster Computing and the Grid*, 2006.

This case study describes FBI Major Case 216, which ultimately became a collaborative investigation between the FBI and site security professionals into a series of cyberattacks that took place from August 2003 to March 2005. Incident response specialists at the

Ultimately, the intrusions were traced back to a 19-year-old man in Uppsala, Sweden, nicknamed "Stakkato," who had begun the attacks when he was 16. Convicted of having gained unauthorized access to several Swedish university networks, "Stakkato" is still under investigation by the FBI for the Cisco code theft [1].

More Contemporary Threats…

June 9, 2014

US Researcher Caught Mining for Bitcoins on NSF Iron

Tiffany Trader

The National Science Foundation has banned a researcher for using agency-funded supercomputers to mine bitcoins, a virtual currency that can be converted into traditional currencies through exchange markets. According to a recently surfaced report from the National Science Foundation Office of the Inspector General, the NSF banned the unnamed researcher after receiving reports that NSF systems at two universities had been used for personal gain.

Bitcoin mining refers to how the virtual currency is generated. Miners solve math problems that serve to verify bitcoin transactions. In exchange they are issued a certain number of bitcoins as a reward.

"The researcher misused over $150,000 in NSF-supported computer usage at two universities to generate bitcoins valued between $8,000 and $10,000," according to the March 2014 Semi Annual Report to Congress. "Both universities determined that this was an unauthorized use of their IT systems. The researcher asserted that he was conducting tests on the computers, but neither university had authorized him to conduct such tests — both university reports noted that the researcher accessed the computer systems remotely and may have taken steps to conceal his activities, including accessing one supercomputer through a mirror site in Europe."

February 12, 2018

Nuclear scientists working at the All-Russian Research Institute of Experimental Physics (RFNC-VNIIEF) have been arrested for using lab supercomputing resources to mine crypto-currency, according to a report in Russia's Interfax News Agency. Located at the Federal Nuclear Center in the Russian city of Sarov, the site is home to a 1 petaflops (peak) supercomputer, installed in 2011. Due to the organization's high secrecy level the supercomputer is not publicly ranked, although it's purported to have a Linpack score of 780 teraflops.

The scientists' plans were foiled when they attempted to connect the classified nuclear resource to the internet.

HPC Has Many of the Same Challenges as Ordinary IT Systems

..and the thousands of probes, scans, stolen credentials, brute-force login attempts, and exploit attempts against hardware, software, and configuration vulnerabilities in HPC facilities today.

What should we do (and not do) about these threats?

We've been thinking about this for a while…

DOE Cybersecurity R&D Challenges for Open Science:

Developing a Roadmap and Vision

American Geophysical Union Building (AGU)

Washington DC

January 24-26, 2007

Meeting Organizers: Deb Agarwal (LBNL), Walter Dykas (ORNL) , and Mike Robertson (DOE)

A Scientific Research and Development Approach To Cyber Security

December 2008

Submitted to

The Department of Energy

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

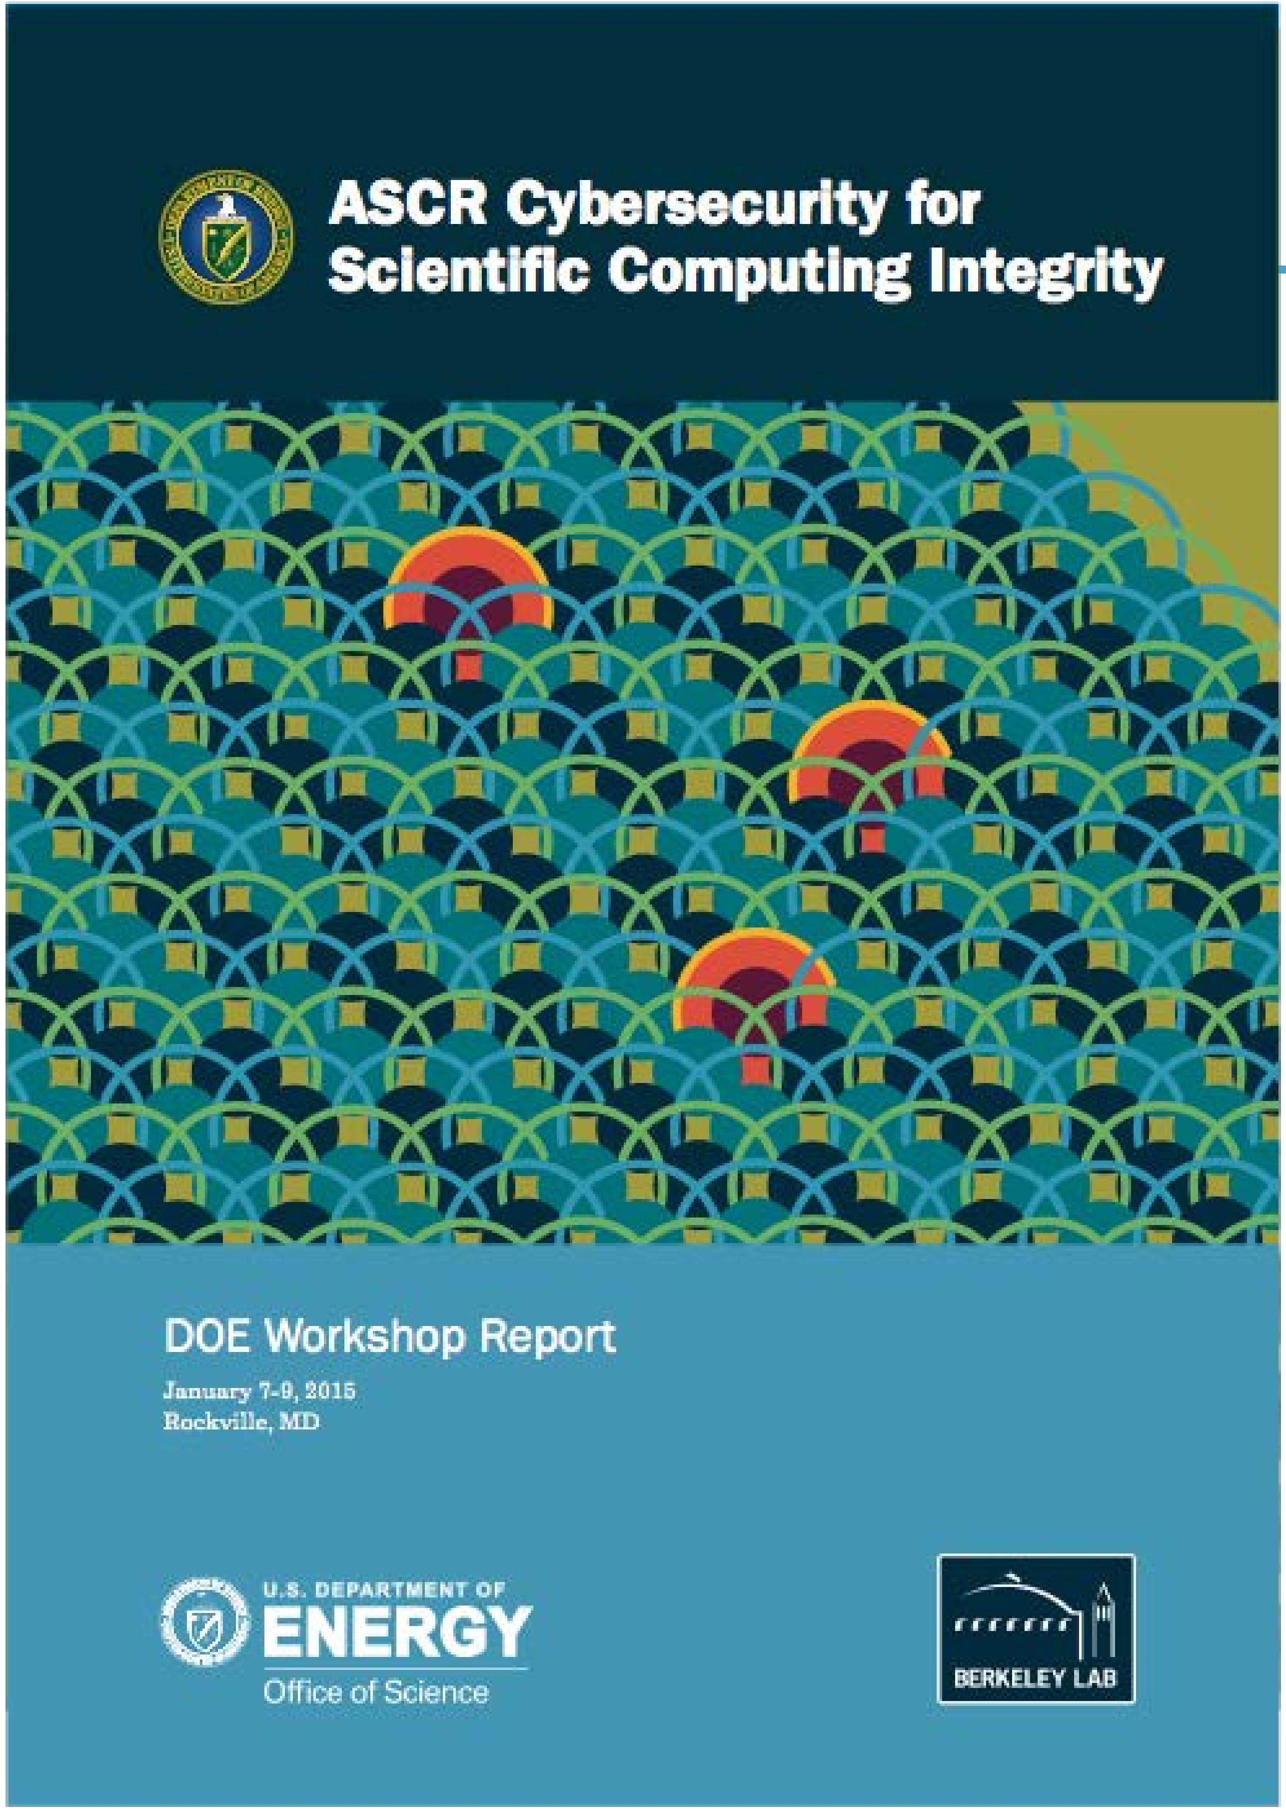

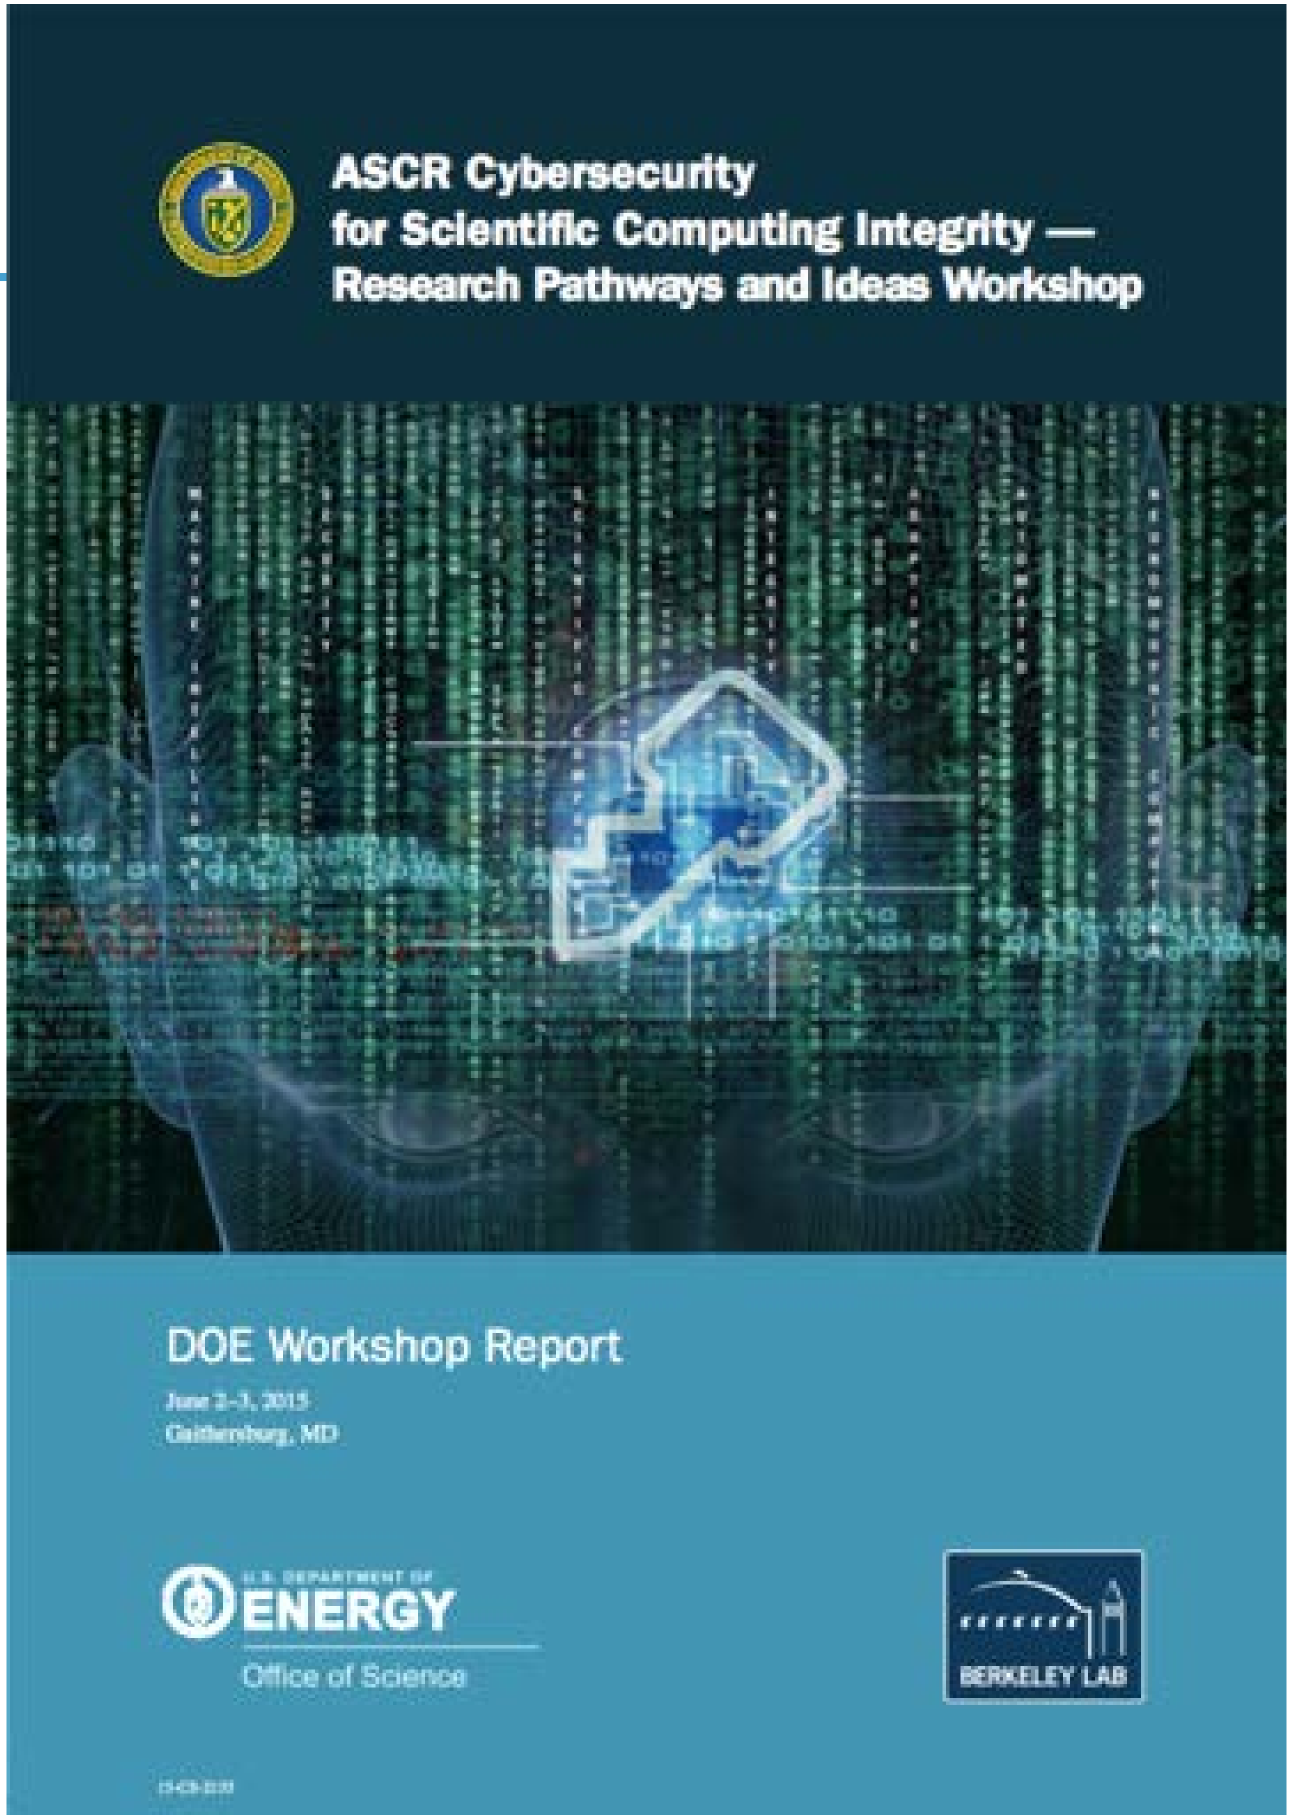

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

Cybersecurity Innovation for Cyberinfrastructure (CICI)

PROGRAM SOLICITATION
NSF 18-547

REPLACES DOCUMENT(S):
NSF 17-528

National Science Foundation

Directorate for Computer & Information Science & Engineering
Office of Advanced Cyberinfrastructure

Full Proposal Deadline(s) (due by 5 p.m. submitter's local time):

June 04, 2018

IMPORTANT INFORMATION AND REVISION NOTES

This solicitation updates the Cybersecurity Innovation for Cyberinfrastructure (CICI) solicitation NSF 17-528. The CICI program continues to support the goal of a secure scientific workflow. The current solicitation:

- Adds two new program areas, Collaborative Security Response Center and Research Data Protection;
- Removes the Cybersecurity Enhancement Area; and
- Renames the Resilient Security Architecture for Research Cyberinfrastructure program area to Secure Scientific Cyberinfrastructure.

NSCI: High-Performance Computing Security Workshop

In July of 2015, the President of the United States issued Executive Order 13702 to create a National Strategic Computing Initiative (NSCI). The goal of the NSCI is to maximize the benefits of High-Performance Computing (HPC) for economic competitiveness and scientific discovery. Security for HPC systems is essential for HPC systems to provide the anticipated benefits. The purpose of this workshop is to identify security priorities and principles that should be incorporated into the strategy of the NSCI, to bring together stakeholders from industry, academia, and Government, and also to identify gaps that should be addressed.

WORKSHOP

September 29, 2016 to Se
30, 2016

NIST 100 Bureau Drive;

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

COMMUNICATIONS OF THE ACM | SEPTEMBER 2017 | VOL. 60 | NO. 9

DOI:10.1145/3096742

Exploring the many distinctive elements that make securing HPC systems much different than securing traditional systems.

BY SEAN PEISERT

Security in High-Performance Computing Environments

Sean Peisert, *Communications of the ACM,* 60(9:72–80), Sept. 2017.

High-Performance Computing Security Workshop

On July 2015, the National Strategic Computing Initiative (NSCI) was established to maximize the benefits of High-Performance Computing (HPC) for economic competitiveness and scientific discovery. For HPC systems to deliver their anticipated benefits, their security requirements must be adequately addressed. To that effect, NIST hosted a workshop in September 2016 that brought together stakeholders from industry, academia, and government to gather their perspectives on the state of technology and future directions. As part of that continuing mission, NIST will host a workshop on March 27-28, 2018 to: review progress in this technology area; assess the threat environment based upon findings and field experience; and build a foundation for development of consensus security principles and controls that are appropriate for the HPC ecosystem.

WORKSHOP

March 27, 2018 t

NIST, 100 Bureau
Gaithersburg, M
Auditorium

What have we learned over the years?

HPC and Traditional IT: Similarities

- Similarities
  - On the surface…
    - Connected to IP networks
    - Often Linux-like OS
      - Similar hardware, software, & configuration challenges and flaws as other systems

HPC and Traditional IT: Differences

- High performance!
  - Computation
  - Data transfers

- Also, many HPC systems (NSF, DOE ASCR) are extremely open, including international collaborations.
  - Can't just "air gap" the HPC system.

∴ ***Can't use certain security solutions, such as network firewalls in the same way***

- Security that impedes collaboration or reduces usability hinders science.

- Some solutions exist that can help compensate for these constraints

Need security without the more simple (but heavy-handed) approaches such as firewalls and air-gaps.

Throughput vs. increasing latency on a 10Gb/s link with *0.0046%* packet loss

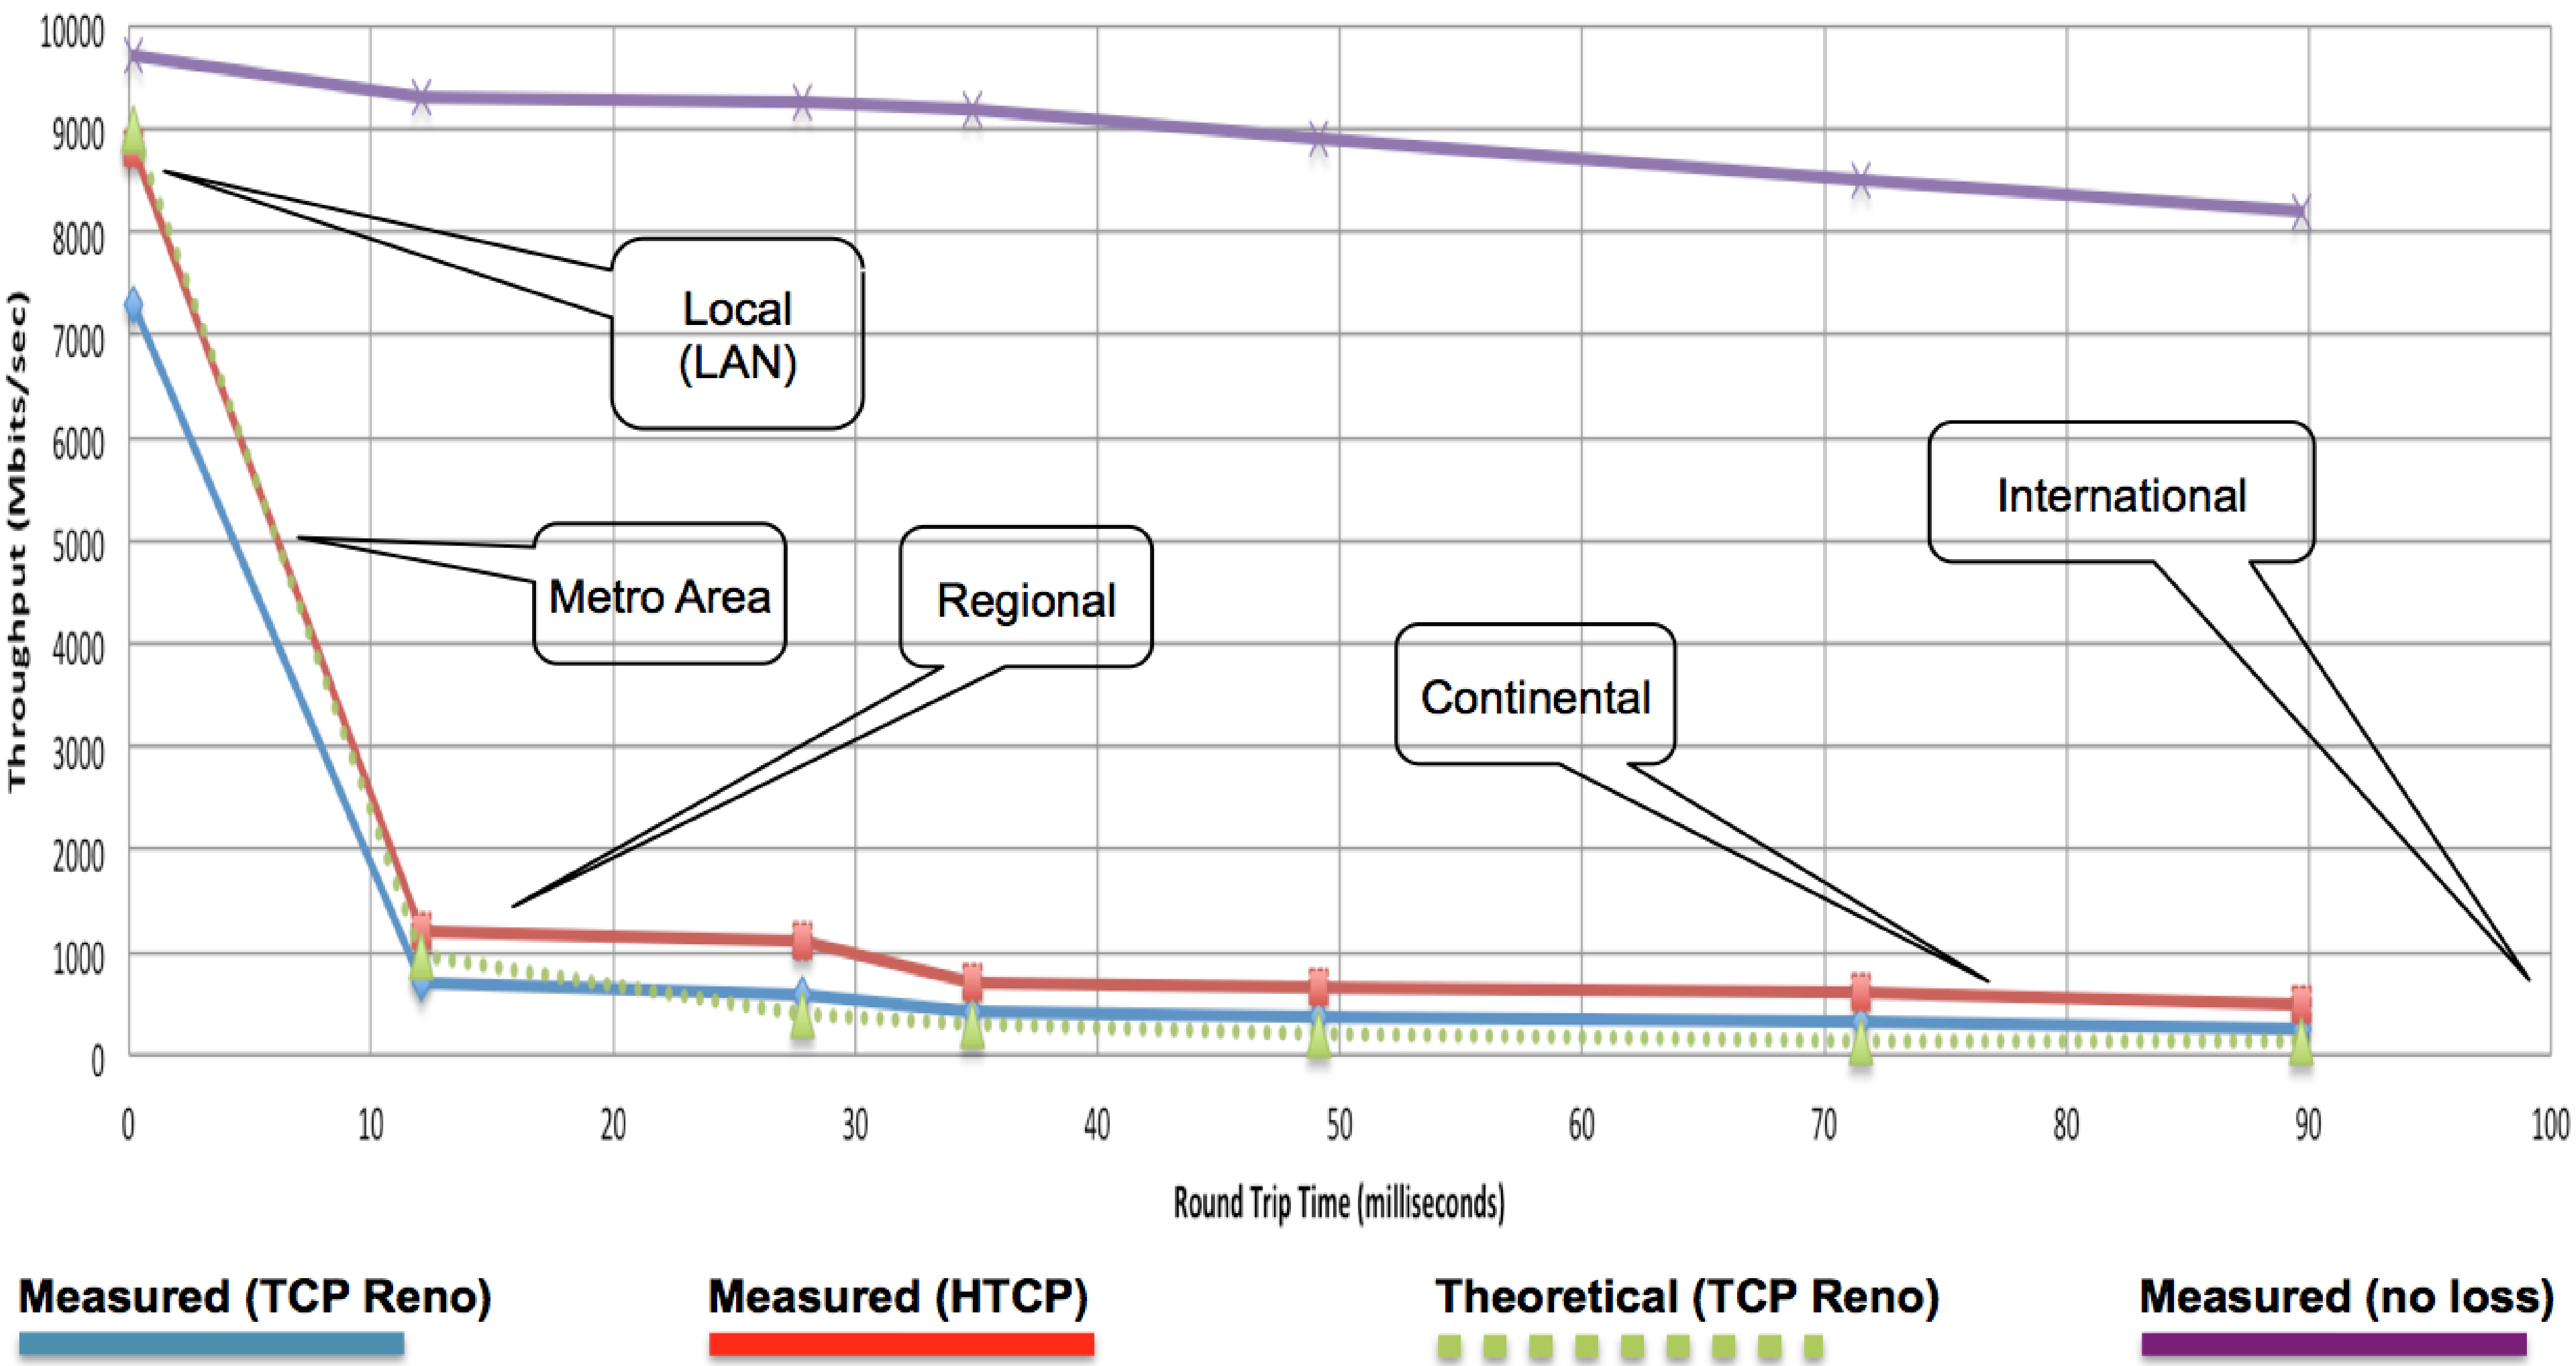

Source: https://fasterdata.es.net/network-tuning/tcp-issues-explained/packet-loss/

Some solutions exist that can help compensate for these constraints

Science DMZ

- Security model that optimizes network throughput
  - Isolates a site's scientific computing in its own network enclave
  - Directs transfers through ***single network ingress/egress point*** that can be ***monitored*** (e.g., with the ***Bro IDS***) and ***restricted*** (e.g., with router ACLs)
  - Achieves throughput by ***reducing complexity***

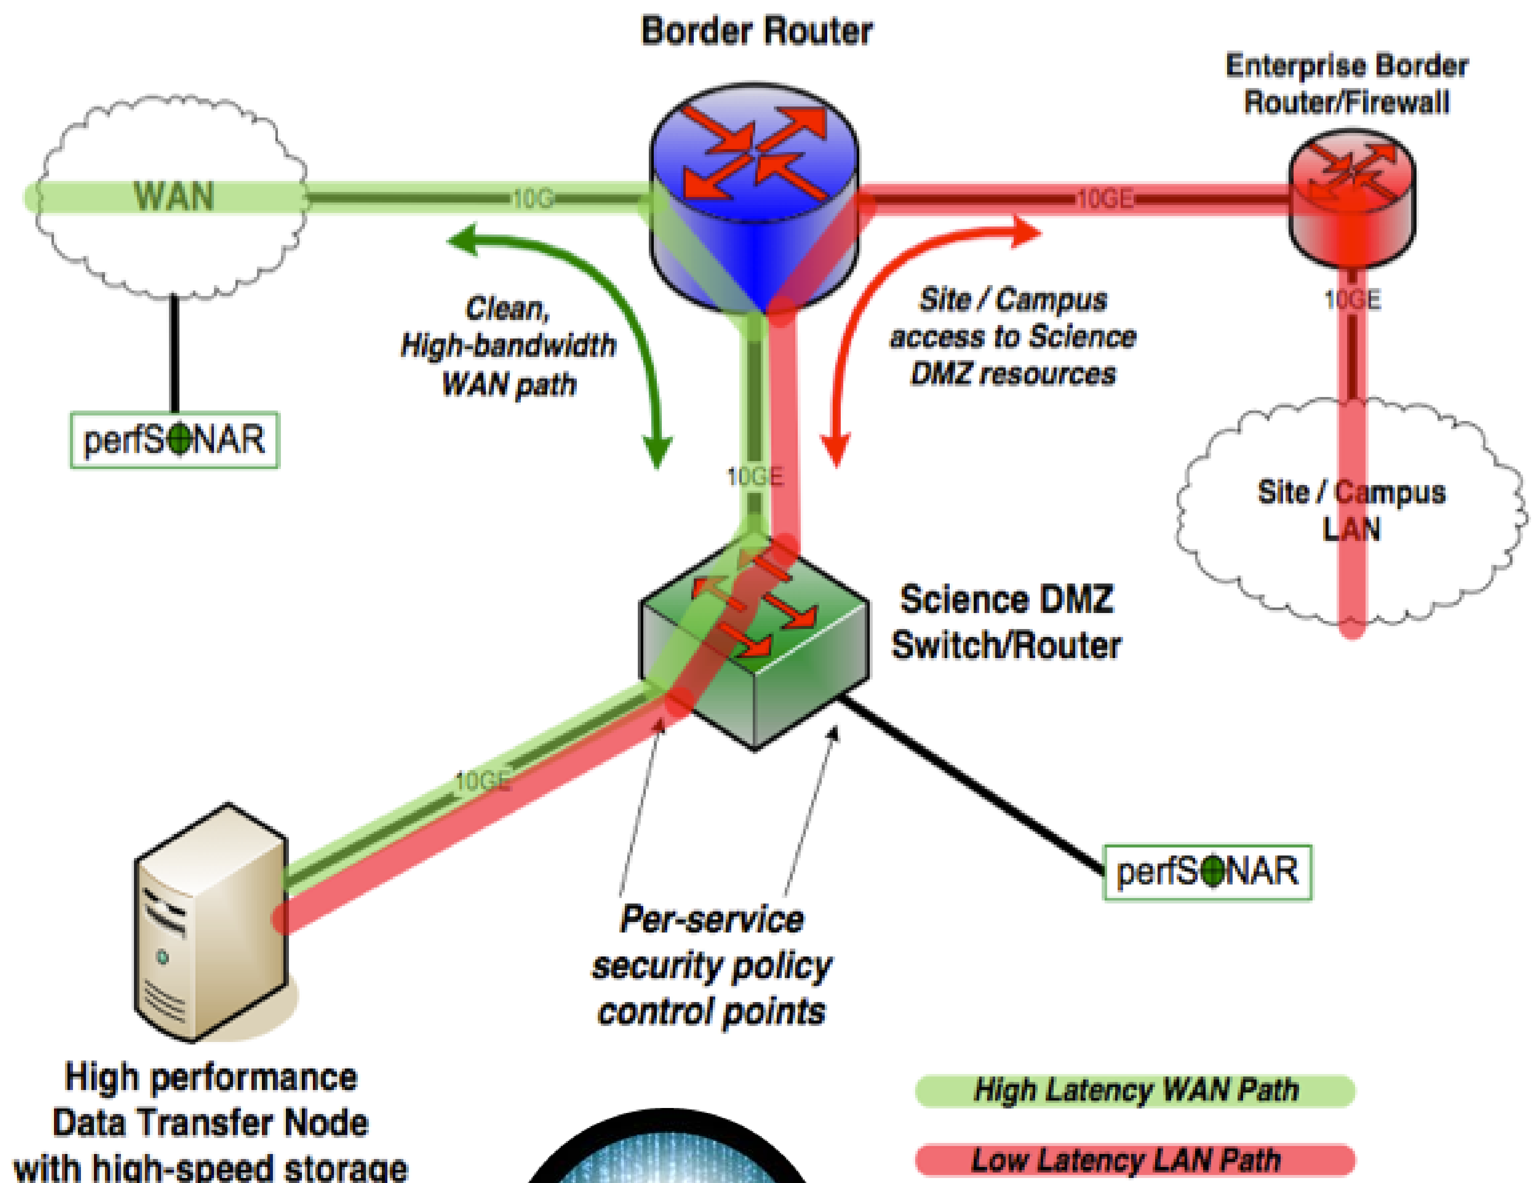

Medical Science DMZ

Journal of the American Medical Informatics Association, 0(0), 2017, 1–8
doi: 10.1093/jamia/ocx104
Research and Applications

AMIA
INFORMATICS PROFESSIONALS. LEADING THE WAY.

OXFORD

Research and Applications

The medical science DMZ: a network design pattern for data-intensive medical science

Sean Peisert,[1,2,3] Eli Dart,[4] William Barnett,[5] Edward Balas,[6] James Cuff,[7] Robert L Grossman,[8] Ari Berman,[9] Anurag Shankar,[10] and Brian Tierney[4]

- Applies Science DMZ framework to computing environments requiring compliance with HIPAA Security Rule

- Key architectures:
  - All traffic from outside compute/storage infrastructure passes through heavily monitored “head nodes.”
  - Storage/compute nodes are not connected directly to the Internet.
  - Traffic containing sensitive or controlled access data is encrypted.

Software engineering is a key goal of the NSCI

IDEAS productivity

Software Engineering for Computational Science and Engineering on Supercomputers

A Birds of a Feather session at SC15, on Wednesday 18 November 2015

- *Robust software can help mitigate vulnerabilities*

The Science DMZ helps compensate for HPC’s limitations — we need more such solutions.

We also need solutions that can leverage HPC distinctiveness as a strength — think different(ly).

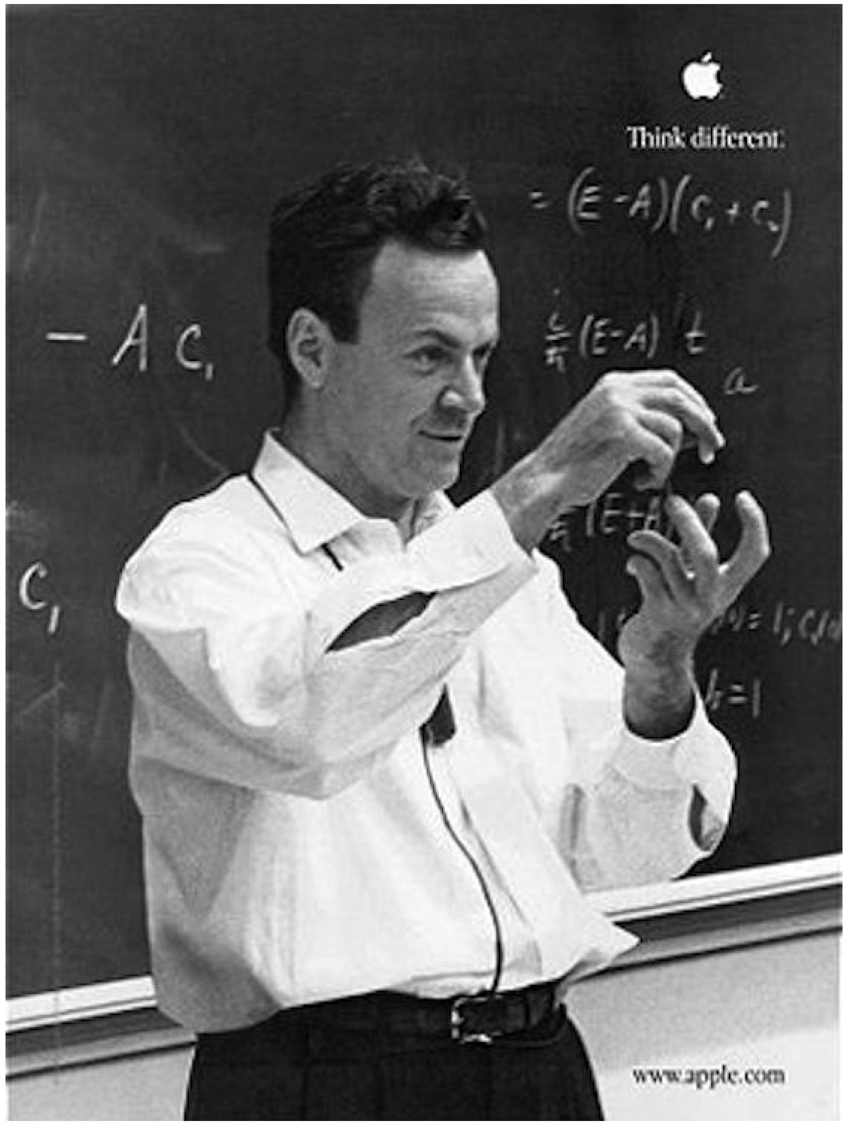

Fingerprinting Computation on HPC Systems

- What are people running on HPC systems?
  - Are they running what they usually run?
  - Are they running what they requested cycle allocations to run?
  - Are they running something illegal (e.g., classified?)

June 9, 2014

US Researcher Caught Mining for Bitcoins on NSF Iron

Tiffany Trader

The National Science Foundation has banned a researcher for using agency-funded supercomputers to mine bitcoins, a virtual currency that can be converted into traditional currencies through exchange markets. According to a recently surfaced report from the National Science Foundation Office of the Inspector General, the NSF banned the unnamed researcher after receiving reports that NSF systems at two universities had been used for personal gain.

Intrusion Detection

Outside the Closed World:
On Using Machine Learning For Network Intrusion Detection

Robin Sommer
International Computer Science Institute, and
Lawrence Berkeley National Laboratory

Vern Paxson
International Computer Science Institute, and
University of California, Berkeley

"…machine learning is rarely employed in operational "real world" settings. … task of finding attacks is fundamentally different from … other applications,

"… Network traffic often exhibits much more diversity .. which leads to misconceptions about what anomaly detection … can realistically achieve…"

"… we argue for the importance of … insight into … an anomaly detection system from an operational point of view. It is crucial to acknowledge [the difficulty in making] progress … without any semantic understanding…"

R. Sommer and V. Paxson, "Outside the Closed World: On Using Machine Learning for Network Intrusion Detection," *Proc. 31st IEEE Symposium on Security & Privacy*, May 2010.

What makes security for HPC different?

- HPC systems tend to:
  - have very *distinctive modes of operation*; or
  - be *used for very distinctive purposes*, notably mathematical computations;
- Some HPC systems:
  - run highly *exotic hardware and software stacks*, and/or
  - are *extremely "open"* to users.

- *This distinctiveness presents both **opportunities** and **challenges***

Key Point #1: HPC systems tend to be used for very distinctive purposes, notably mathematical computations

Data Analysis

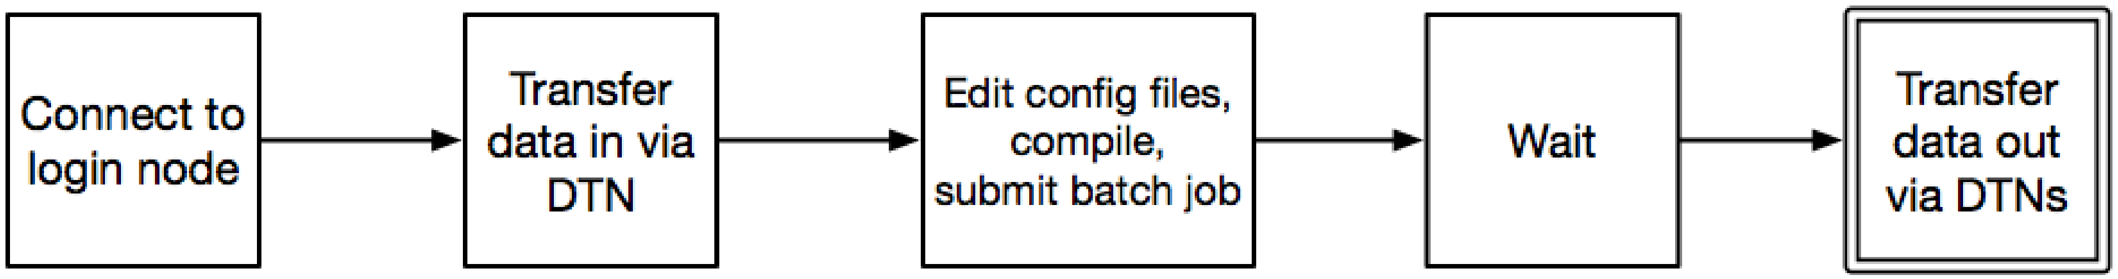

Simulation

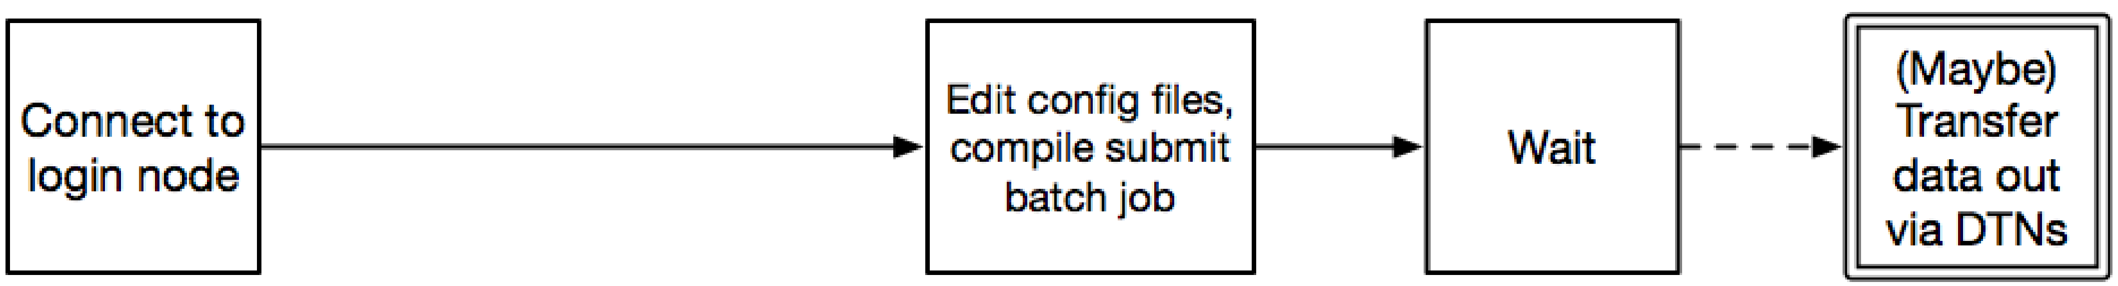

Simulation with Coupled Computation/Visualization

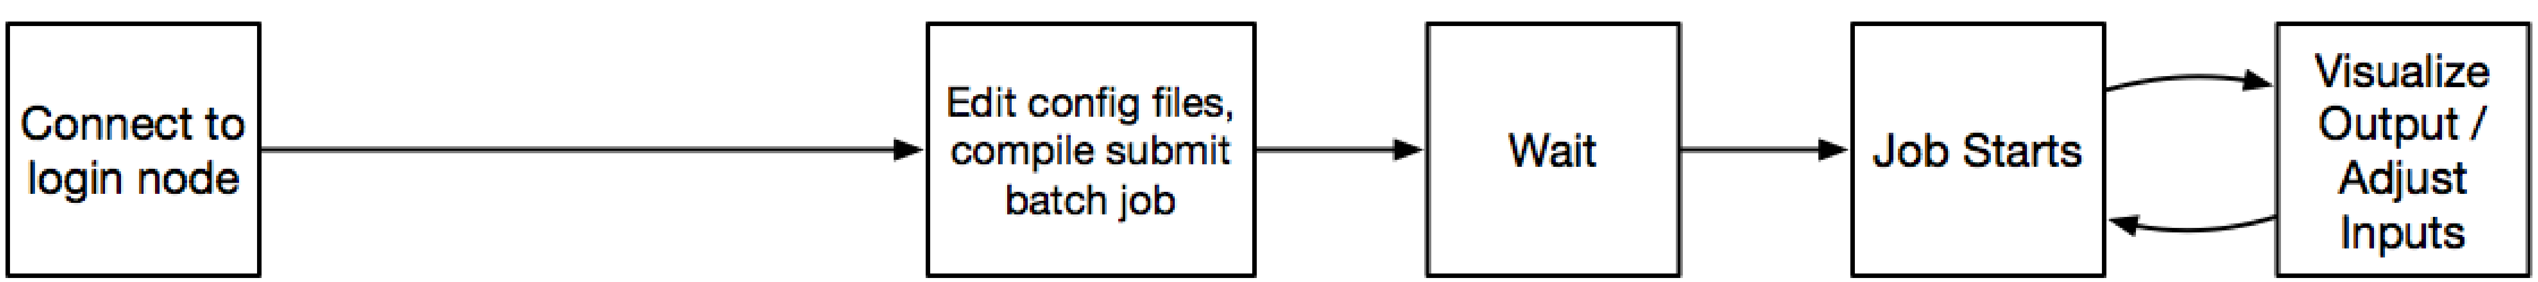

Analytics vs. Simulation Kernels:

7 Giants of Data	*7 Dwarfs of Simulation*
Basic statistics	***Monte Carlo methods***
Generalized N-Body	***Particle methods***
Graph-theory	***Unstructured methods***
Linear algebra	***Dense linear algebra***
Optimizations	***Sparse linear algebra***
Integrations	***Spectral methods***
Alignment	***Structured meshes***

Source: K. Yelick, "A Superfacility for Data Intensive Science," ASCAC Meeting, Sept. 2016.

Key Point #2: What if there was less diversity in the events, and greater semantic understanding?

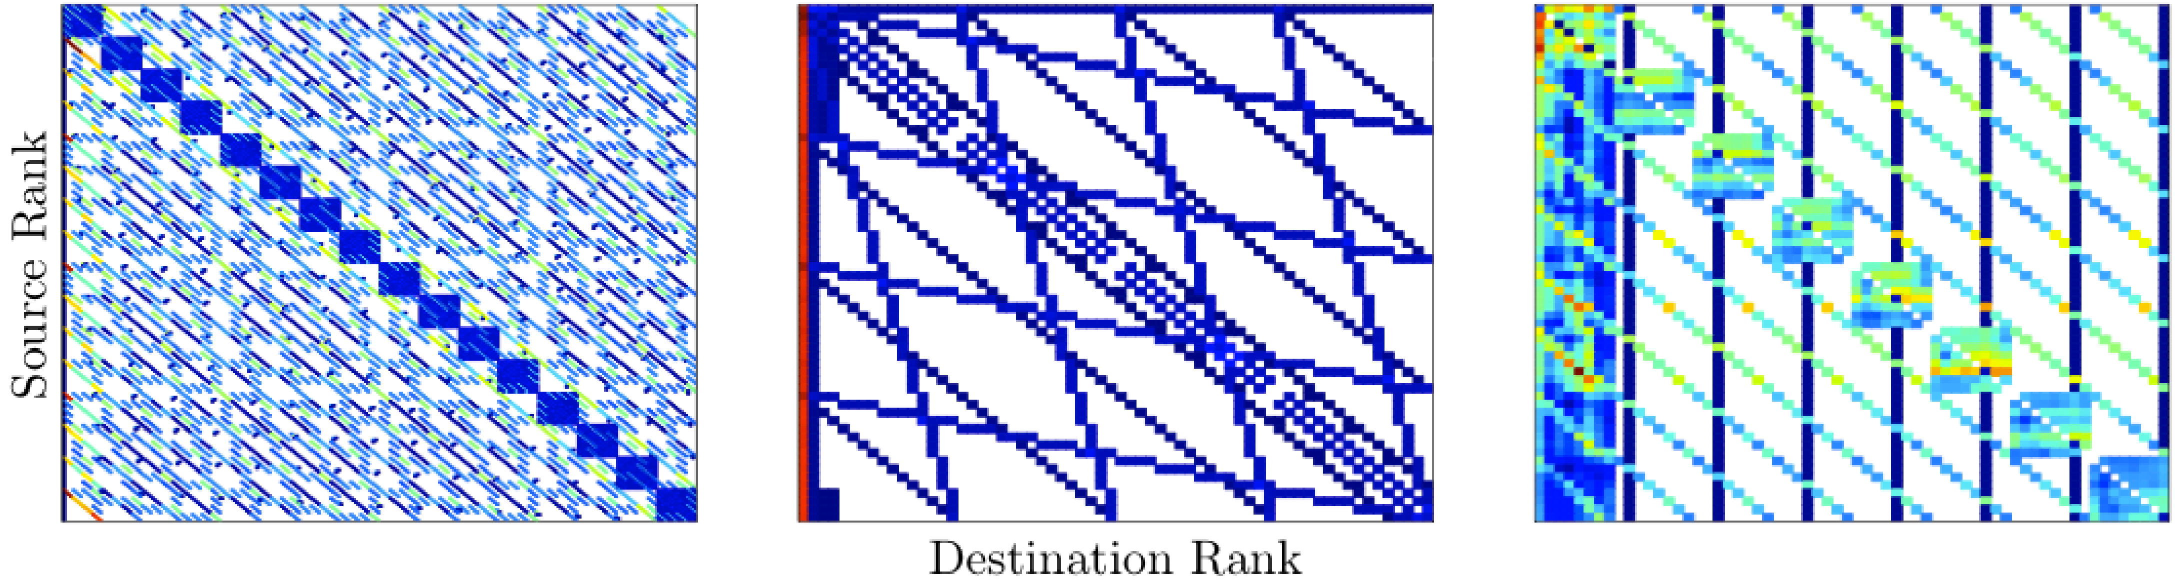

- Developed technique for fingerprinting computation on HPC systems
- Used hundreds of ***MPI logs*** and ***time-series CPU information*** for dozens of scientific applications from NERSC HPC systems.
  - Applied Bayesian machine learning for classification of scientific computations.
  - Approach identifies test HPC codes with 95-99% accuracy.

S. Whalen, S. Engle, S. Peisert, and M. Bishop, “Network-Theoretic Classification of Parallel Computation Patterns,” *International Journal of High Performance Computing Applications*, 26(2):159–169, May 2012.

S. Whalen, S. Peisert, and M. Bishop, “Multiclass Classification of Distributed Memory Parallel Computations,” *Pattern Recognition Letters*, 34(3):322–329, February 2013.

B. Copos and S. Peisert (dissertation advisor), *Modeling Systems Using Side Channel Information,* Ph.D. dissertation, University of California, Davis, 2017.

Looking to the future

Looking forward

- *The threat isn't going away*
- *Science is changing*
  - Sensor data
  - Distributed / streaming data collection
- *Science data is getting to us in new ways, and we have more data to protect.*

HP Security Opportunities: Monitoring Data

- Monitoring data is useful for *security monitoring for abnormal behavior*
  - Misuse of cycles
  - Identifying manipulated programs (malware, etc..)

  - Also useful for *provenance / integrity monitoring*

- But… the ability to successfully perform analysis on monitoring data depends on *availability of useful monitoring data*

- *Key Point #3: custom stacks provide opportunities for instrumenting system hardware or software to capture additional audit/provenance data.*

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

Current focus on provenance in HPC might help provide better monitoring data

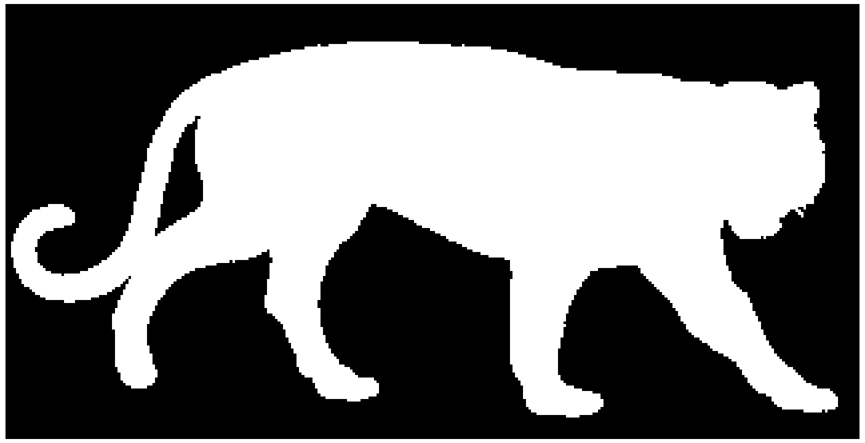

jupyter

KBase
PREDICTIVE BIOLOGY

Provenance

Tracking the user and transformation of data, thus allowing credit to be given to data contributors, analysts, and tool developers in addition to enabling the recording and sharing of methods.

SUMMIT

CNL

GPUs

INTEL® OMNI-PATH FABRIC 100 SERIES

Lustre®

HPC systems that run exotic hardware and software stacks may also provide monitoring data — exascale / quantum / neuromorphic should only continue this

mOS

Mira

Mira Ushers in a New Era of Scientific Supercomputing

As one of the fastest supercomputers, Mira, our 10-petaflops IBM Blue Gene/Q system, is capable of 10 quadrillion calculations per second. With this computing power, Mira can do in one day what it would take an average personal computer 20 years to achieve.

CNK

ECP
EXASCALE COMPUTING PROJECT

NVLINK HIGH-SPEED INTERCONNECT
Designed for Accelerated Computing

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

Trend toward constrained modes of operation

- *Containerization* — all interaction takes place within the container

Trend toward constrained modes of operation

- *Limited interfaces / “Automated Supercomputing”*
  - Science gateways — web portals to HPC
  - “Superfacility” model

Security tends to benefit from more constrained operation, which is the general trend.

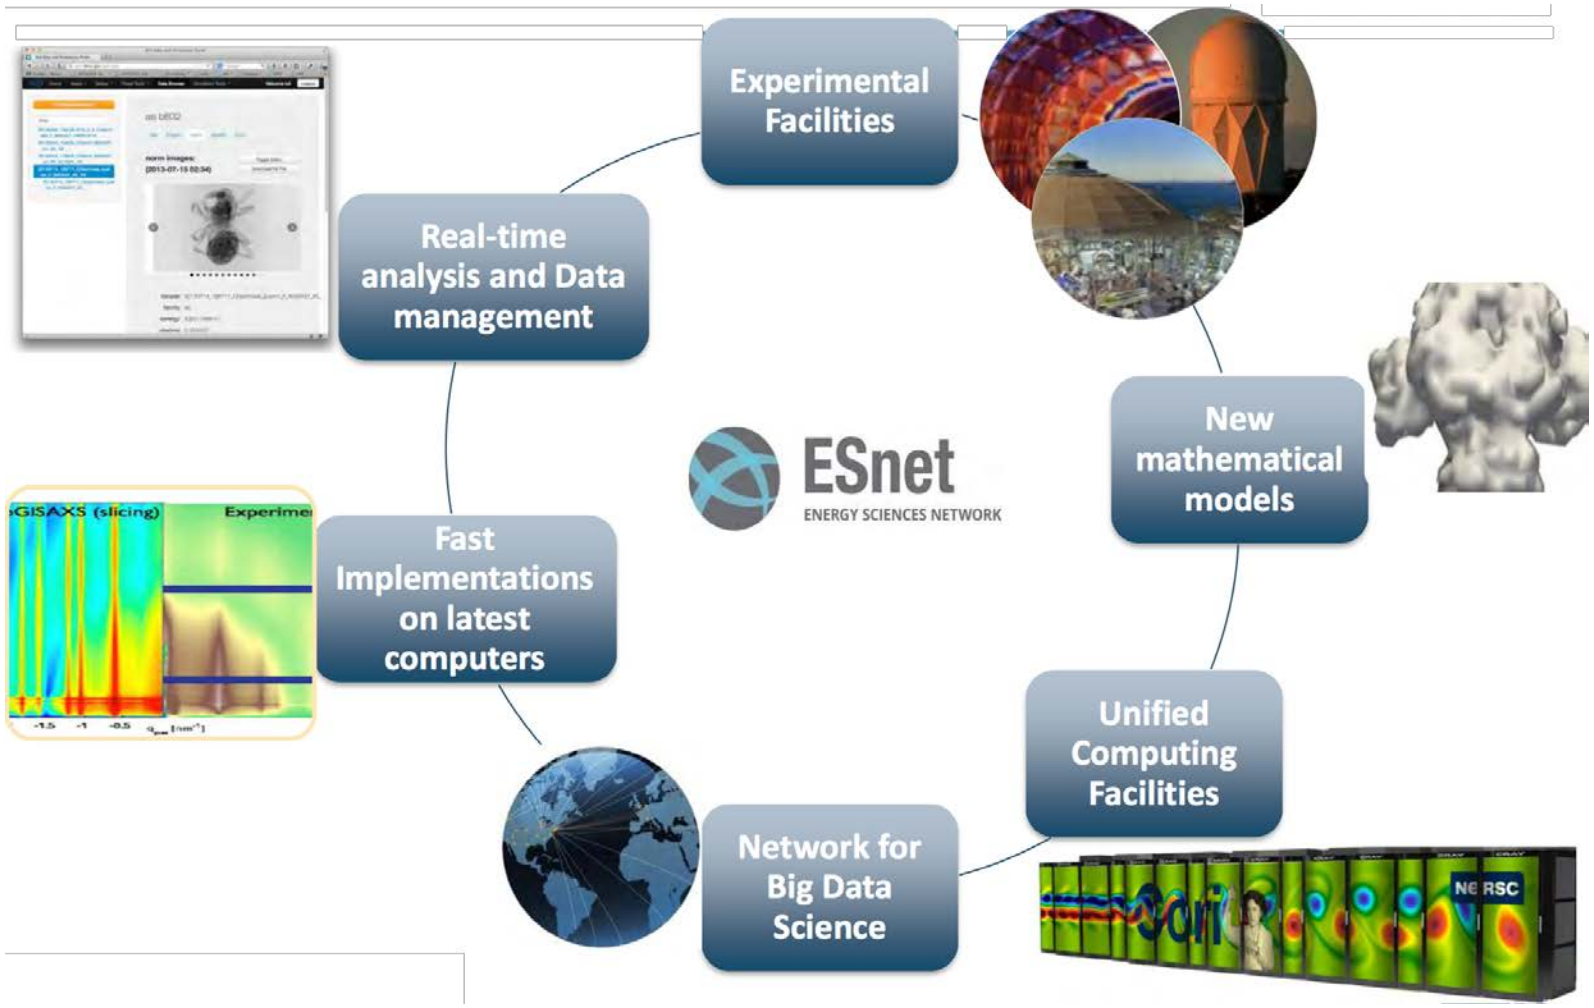

Summary

- HPC systems are different in how they're built and used — *challenges & opportunities*.

- Key security challenges:
  - Traditional security solutions often compete with *priority of high-performance*.
  - Many HPC environments are *highly "open"* to enable broad scientific collaboration.

- Key security opportunities:
  - HPC systems used for *distinctive purposes*, and have strong *"regularity" of activity.*
  - *Custom HW/SW stacks* provide opportunities for enhanced *security monitoring*.
  - Trend toward ***containerized operation*** & ***limited interfaces*** in HPC is likely to help.

My call-to-arms / challenges to you

- ***Make sure*** *you focus security efforts around the most important goals*
  - Data leakage (even in “open science”)
  - Alteration of code or data
  - Misuse of computing cycles
  - Disruption/denial of service against HPC systems or networks

- ***Make sure*** *what you do prioritizes performance and usability / openness*

- ***Think*** *about how …*
  - *we can best influence future HW/SW stack design to provide opportunities for enhanced security monitoring / provenance tracking / etc…*

  - *to accelerate the trend toward* ***containerized operation*** *&* ***limited interfaces****.*

- *Keep an eye on* ***up-and-coming security technologies***
  - e.g., Computing over encrypted data (“somewhat homomorphic encryption”)

Contact:
Dr. Sean Peisert
sppeisert@lbl.gov
http://crd.lbl.gov/Q/peisert/
http://crd.lbl.gov/Q/HPC-Security/

U.S. DEPARTMENT OF ENERGY

UNIVERSITY OF CALIFORNIA

BERKELEY LAB
Lawrence Berkeley National Laboratory

Backup Slides

High-Performance Computing Has Become Essential to U.S. National Security and Prosperity

- *Scientific understanding*
  - cosmology
  - particle physics
  - climate change
  - biological systems
  - renewable energy
  - precision medicine
  - nuclear stockpile safety
- *Engineering analysis*
  - Aerodynamics/hydrodynamics
  - Materials
- *Cryptanalysis*
- ..and more