Training Network Administrators in a Game-Like Environment

Engin Arslan, Murat Yuksel, Mehmet Gunes
enginars@buffalo.edu, yuksem@unr.edu, mgunes@unr.edu

Computer Networking Lab
http://cnl.cse.unr.edu
Computer Science and Engineering
University of Nevada – Reno

NANOG64
June 2005, San Francisco, CA

Acknowledgements

- Collaborators
  - Faculty
    - Mehmet H. Gunes (mgunes@unr.edu), CSE, UNR
    - Ramona Houmanfar (ramonah@unr.edu) , Psychology, UNR
  - Alumnus
    - Engin Arslan (enginars@buffalo.edu) Ph.D.

- Sponsors
  - This work was supported by the U.S. NSF awards 1321069, 0721600, and 0721609.

Motivation

- Network management and automated configuration of large-scale networks is a crucial issue for ISPs
  - SLAs to meet
  - High/strict demand apps: IPTV, VoIP
- Must respond to failures or demand spikes in a timely (24/7) and smart manner
  - Wrong response or configuration translates to real $ costs
- ISPs generally trust experienced administrators to manage network
  - particularly for functions involving dynamic optimization
  - e.g., traffic engineering

Training Network Administrators

- Network administrator training is a long-term process
- Exposing inexperienced administrators to the network is too risky
- Current practice to train is apprenticeship

Can we train the network administrators using a game-like environment rather than months of years of apprenticeship?

IGP Link Weight Setting

- How to set the weights?
  - Inversely proportional to link capacity?
  - Proportional to propagation delay?
  - Network-wide optimization based on traffic?

Dynamism → routing instability

How about scale? > 10K links

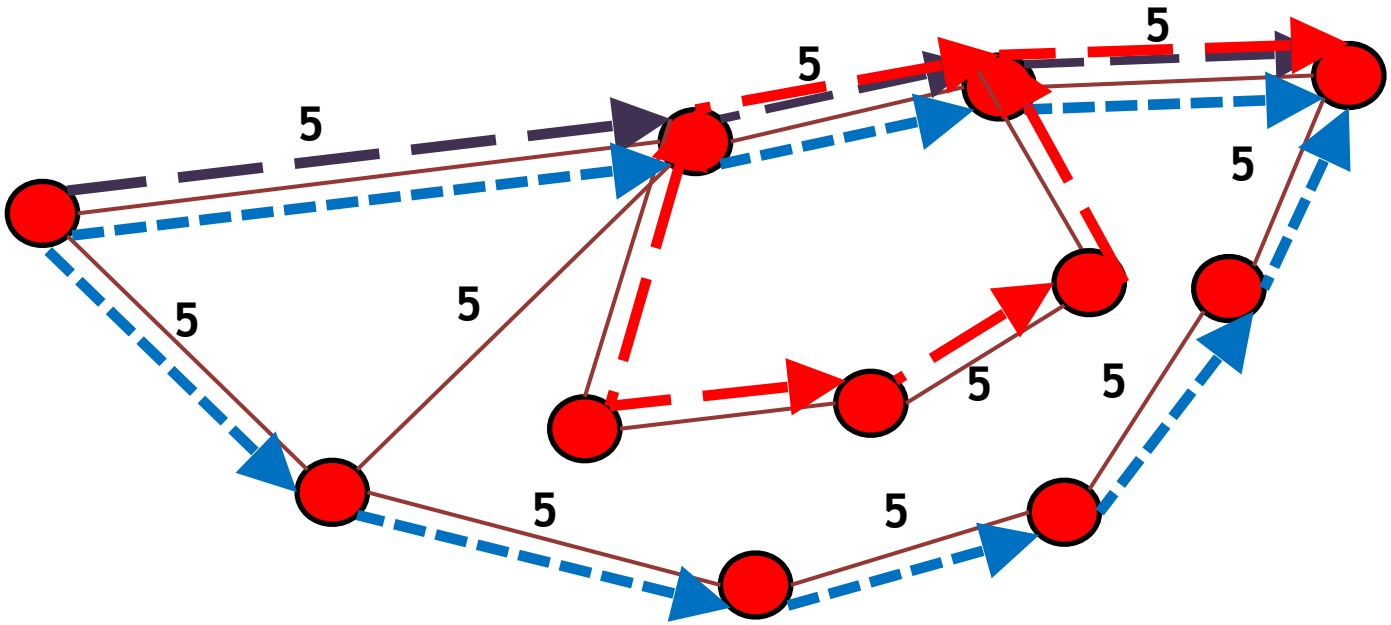

J. Rexford et al., http://www.cs.princeton.edu/courses/archive/spr06/cos461/

IGP Link Weight Setting

- Empirical way:
  - Network administrator experience
  - Trial and error
  - error-prone, not scalable

Routing instabilities scare ISPs

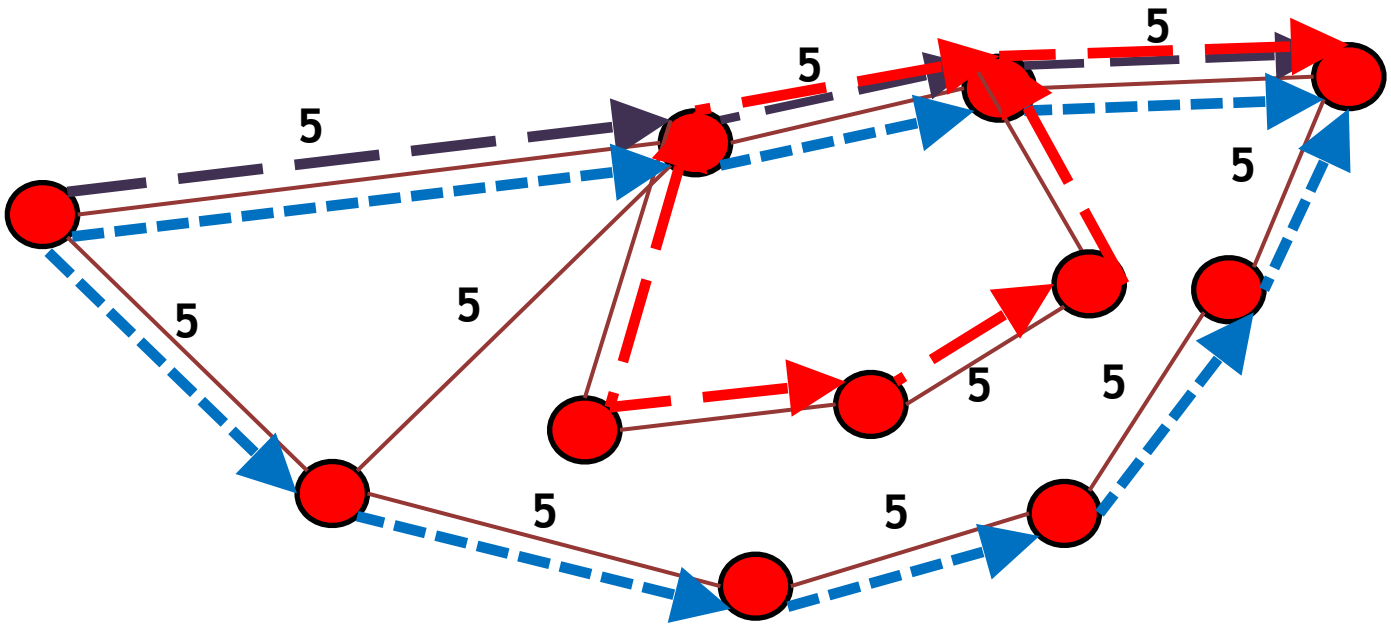

J. Rexford et al., http://www.cs.princeton.edu/courses/archive/spr06/cos461/

Traffic Engineering: IGP Link Weight Setting Problem

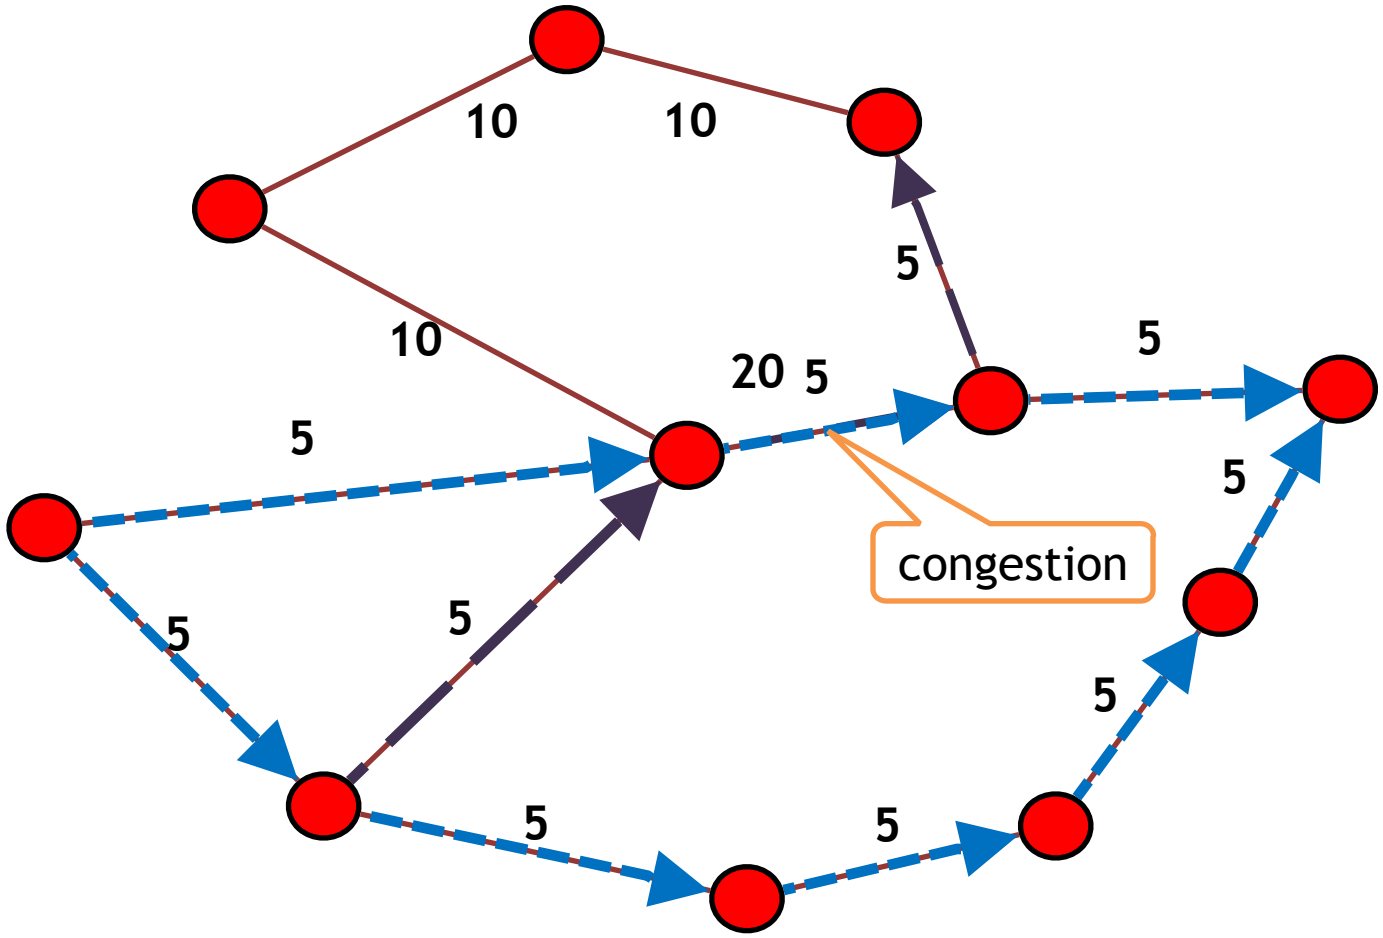

Network Management Game (NMG) Framework

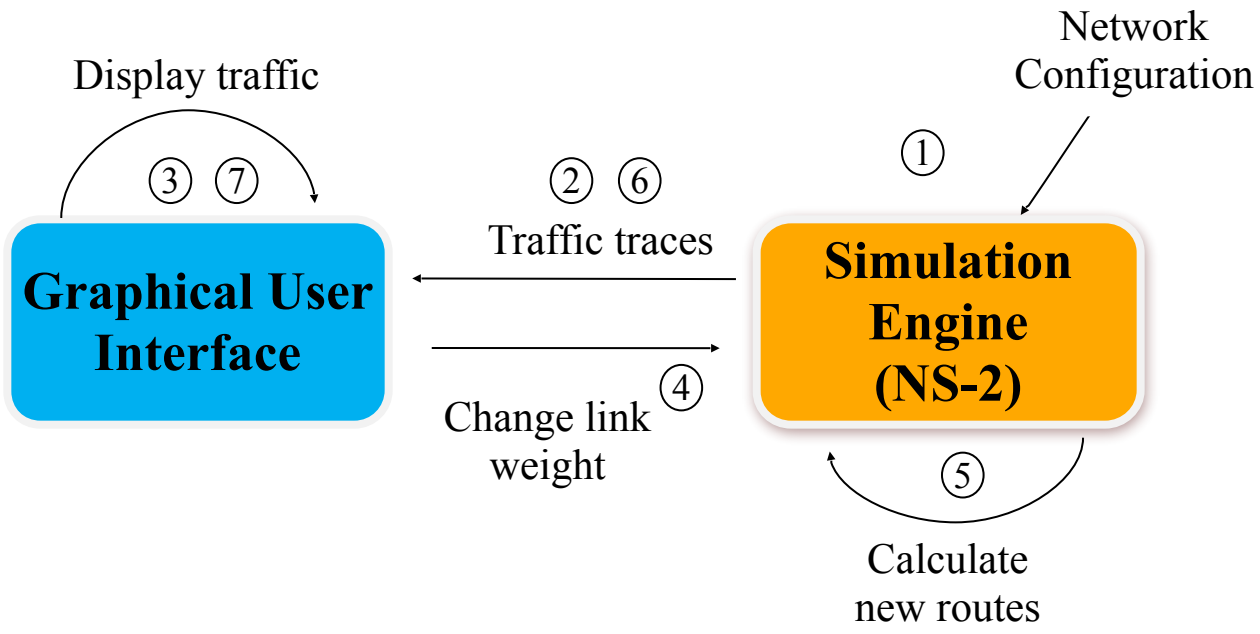

Block diagram of Network Management Game (NMG) components.

Network Simulator (NS-2)

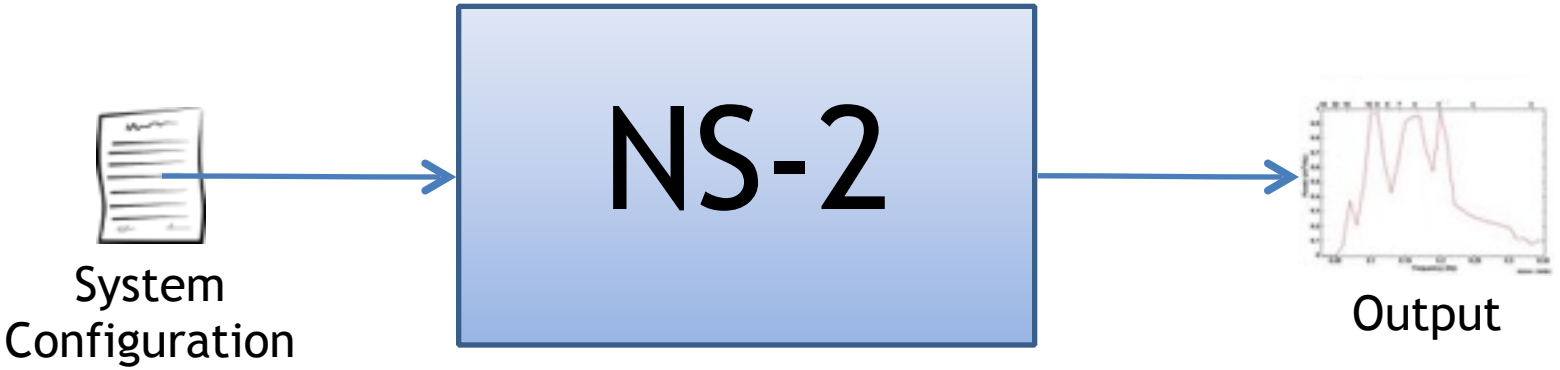

- No real time interactivity

  Run simulation→ See the results
- Necessitates adequate level of TCL scripting
- Not designed for training purpose

Simulator-GUI Interaction

- Concurrency is challenging
  - ❖ Run the simulation engine for a time period then animate in GUI before the engine continues
  - ❖ Slowdown animator – chose this approach
- GUI-Engine interaction is achieved via TCP port
  - ❖ Animator opens a socket to send simulation traces
  - ❖ GUI opens a socket to send commands

Sample Message: *$ns $n1 $n2 2* →

set weight of link between n1 and n2 to 2

NMG Screenshot

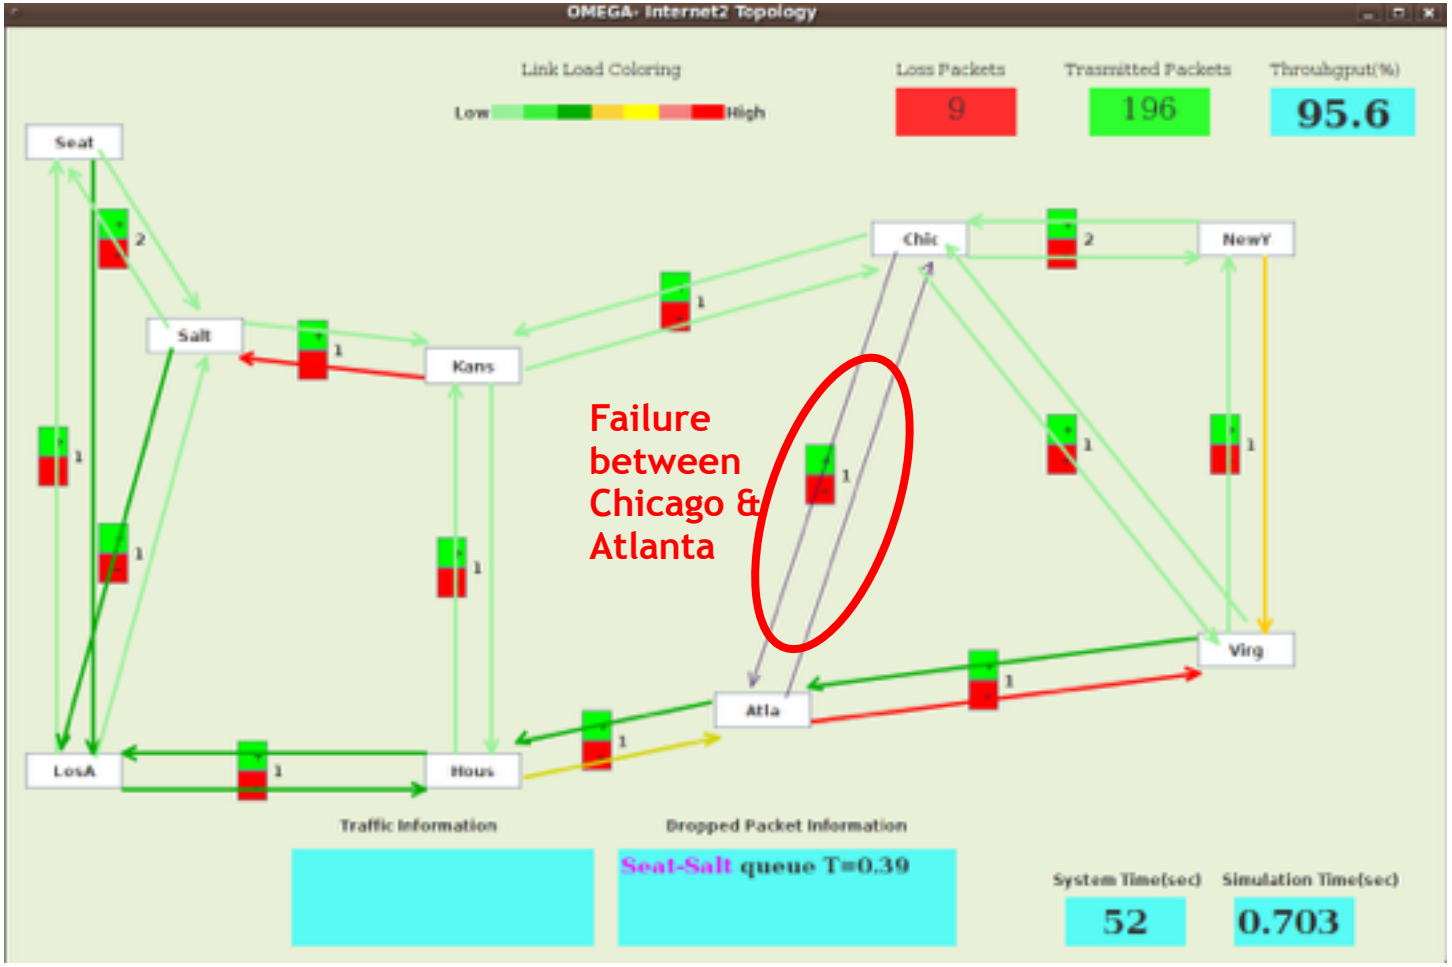

NMG Screenshot

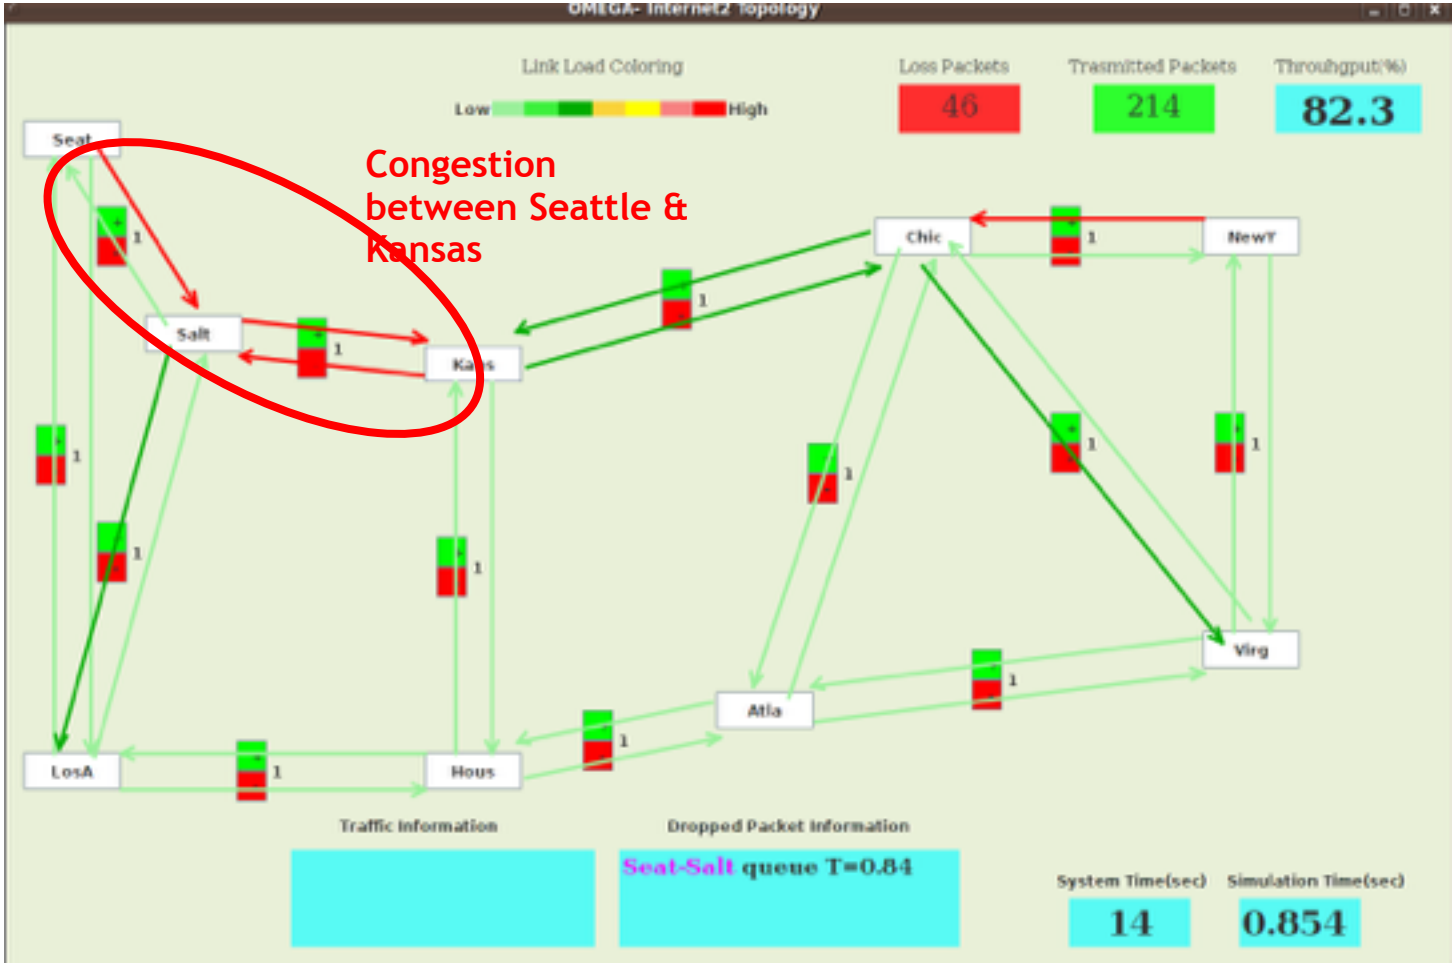

User Goal

- Increase Overall Throughput by manipulating link weights within a given time period

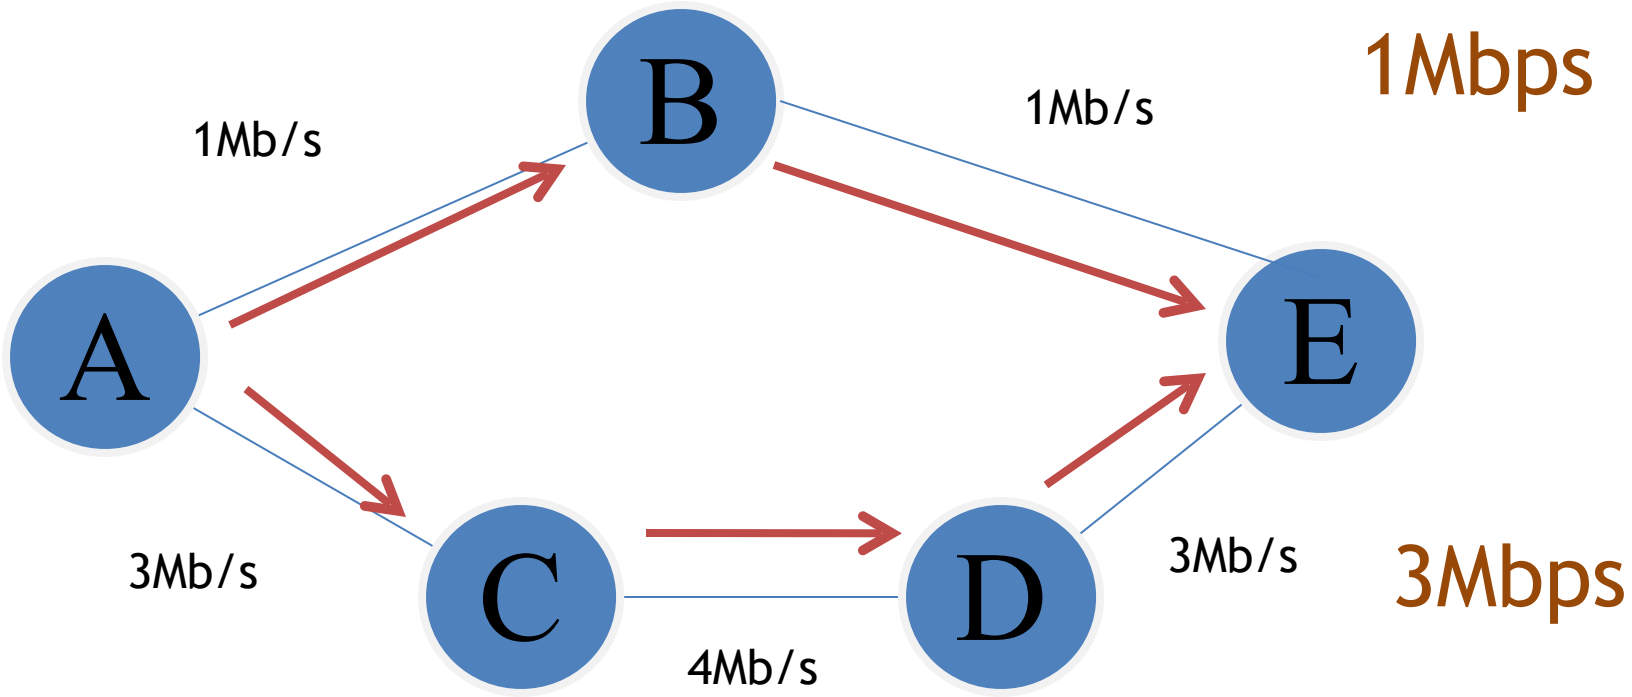

User Experiments

We conducted 2 user experiments

- Training without Mastery
  - ❖ No specific skills targeted
  - ❖ No success level obligated
- Training with Mastery
  - ❖ Two skills are targeted to train
  - ❖ Success level obligated

Life of An Experiment

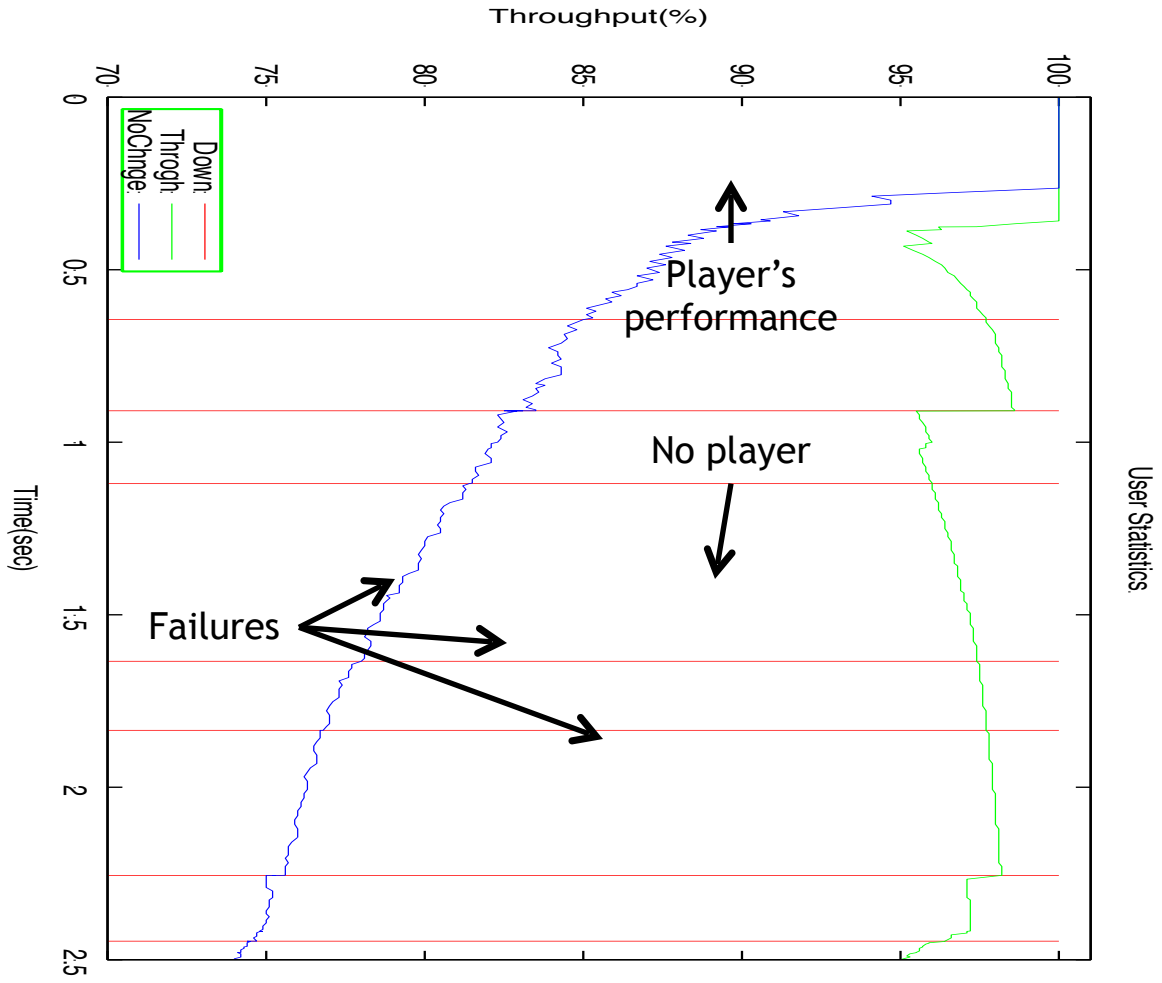

Training without Mastery

- 5 training scenarios
- For every scenario, user has fixed 3-5 minutes to maximize overall throughput
- 8 users attended
- Took around 45 minutes for each user
- User performance evaluated for failure and no failure cases

User Experiment Design

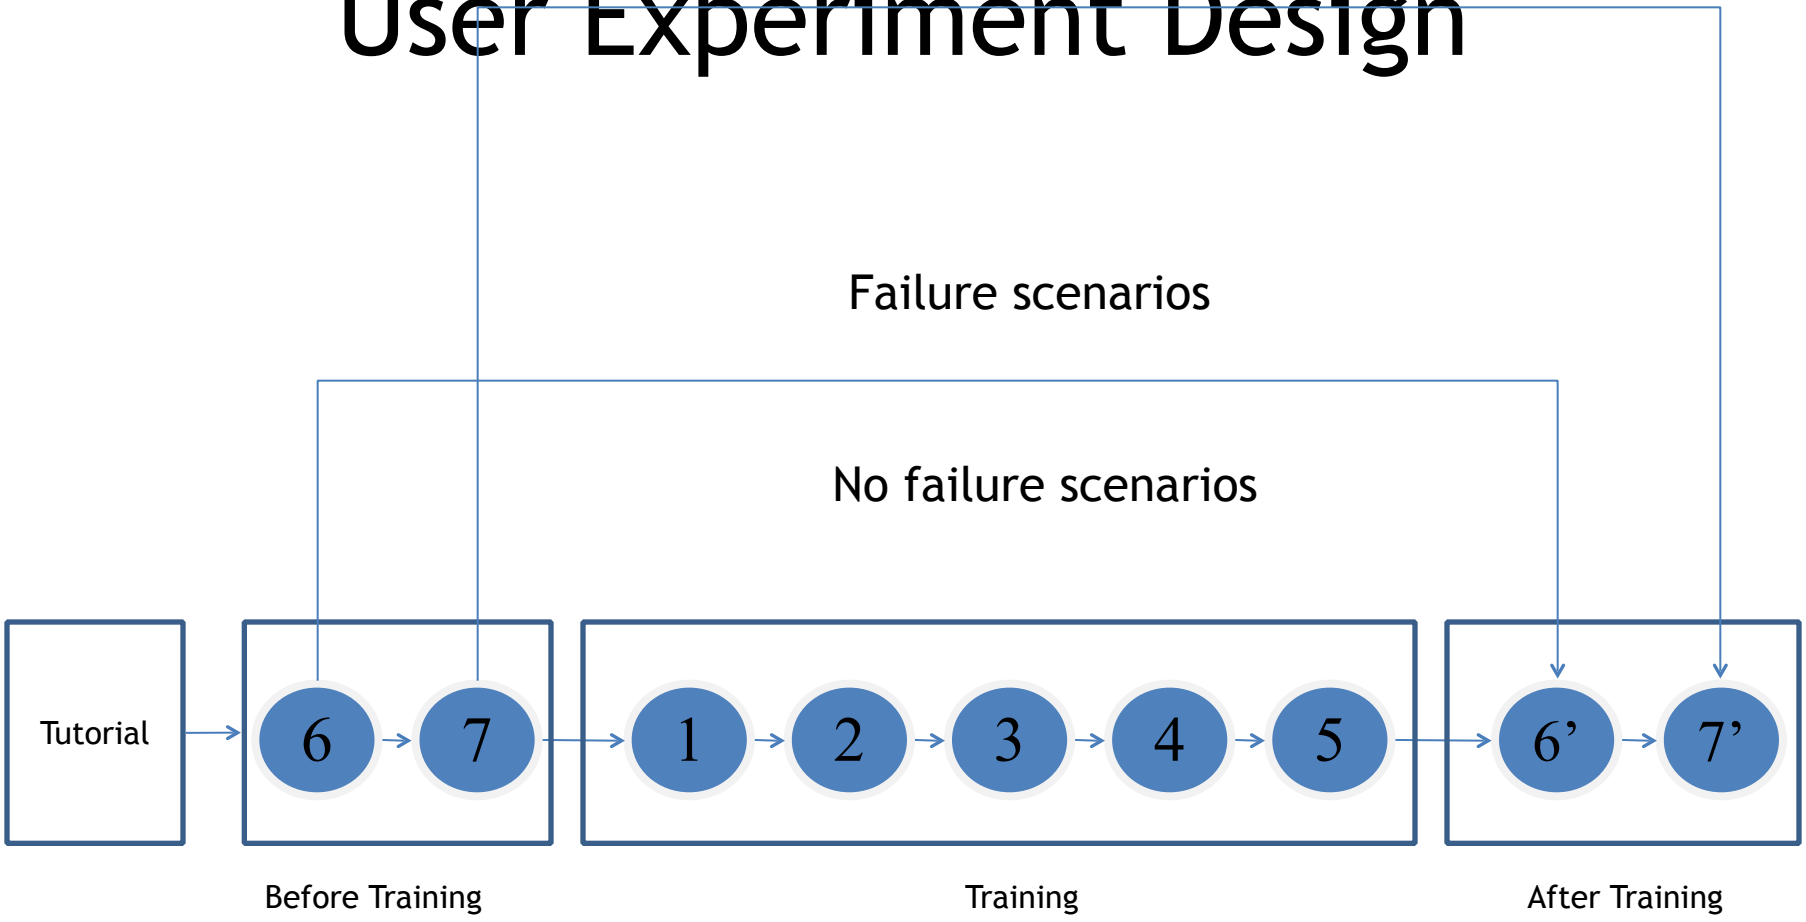

Best and Worst Players

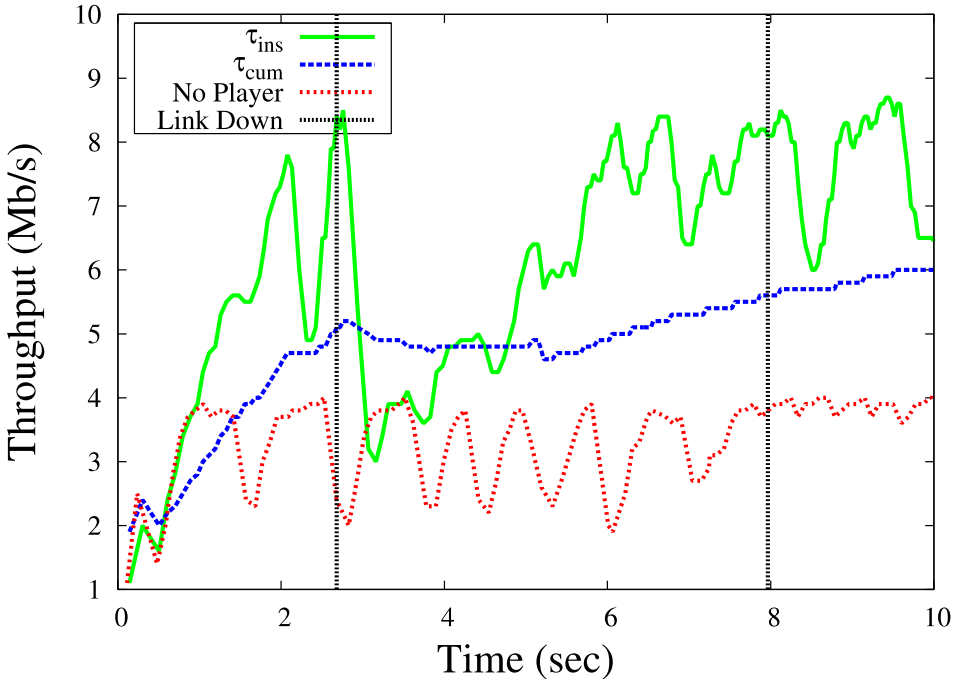

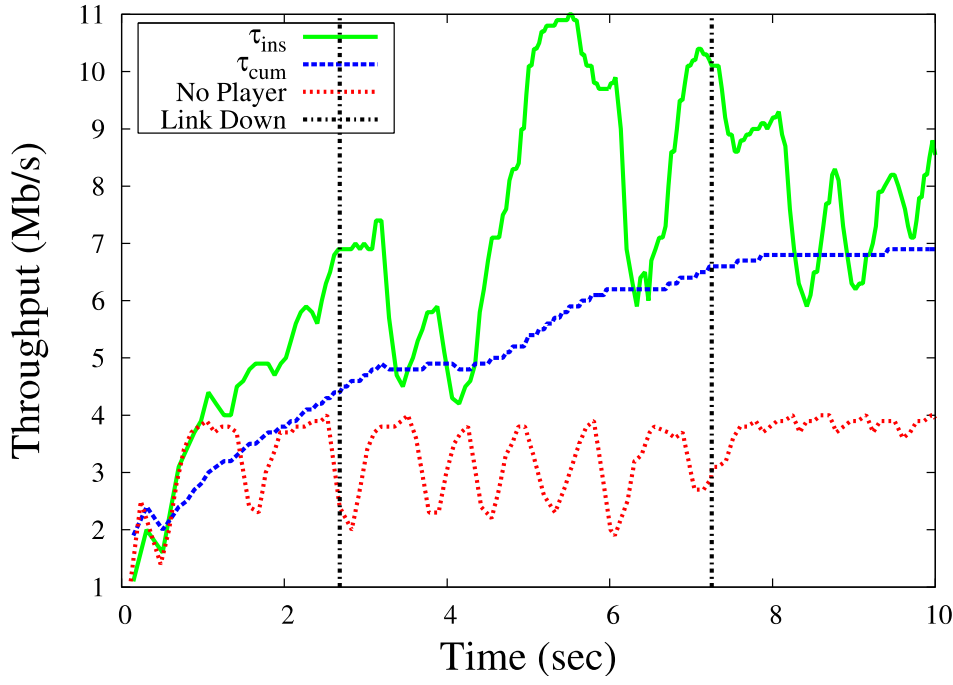

Before Training
Case 7

Best and Worst Players

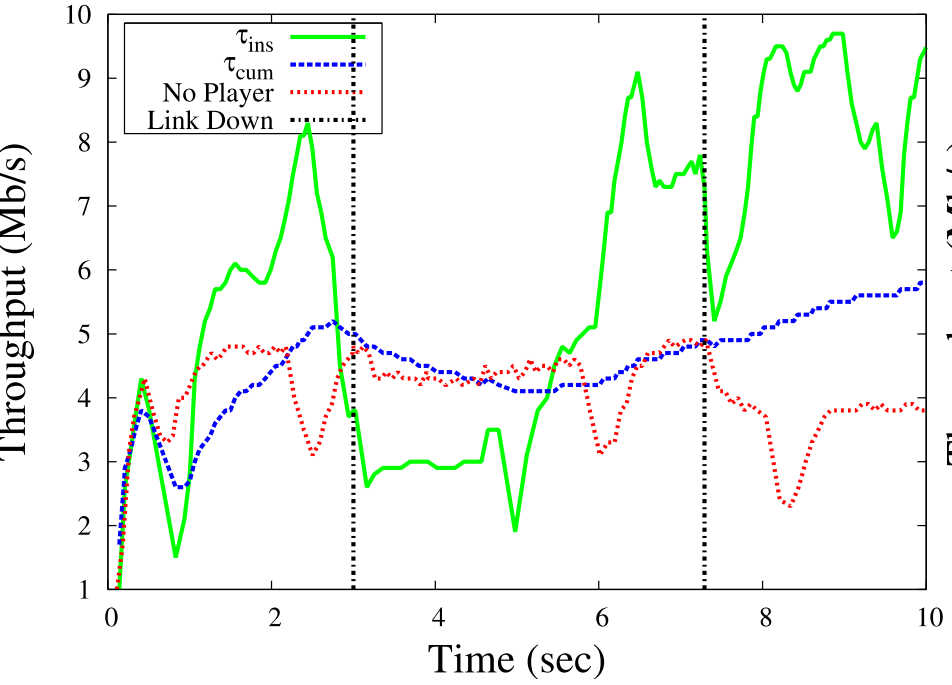

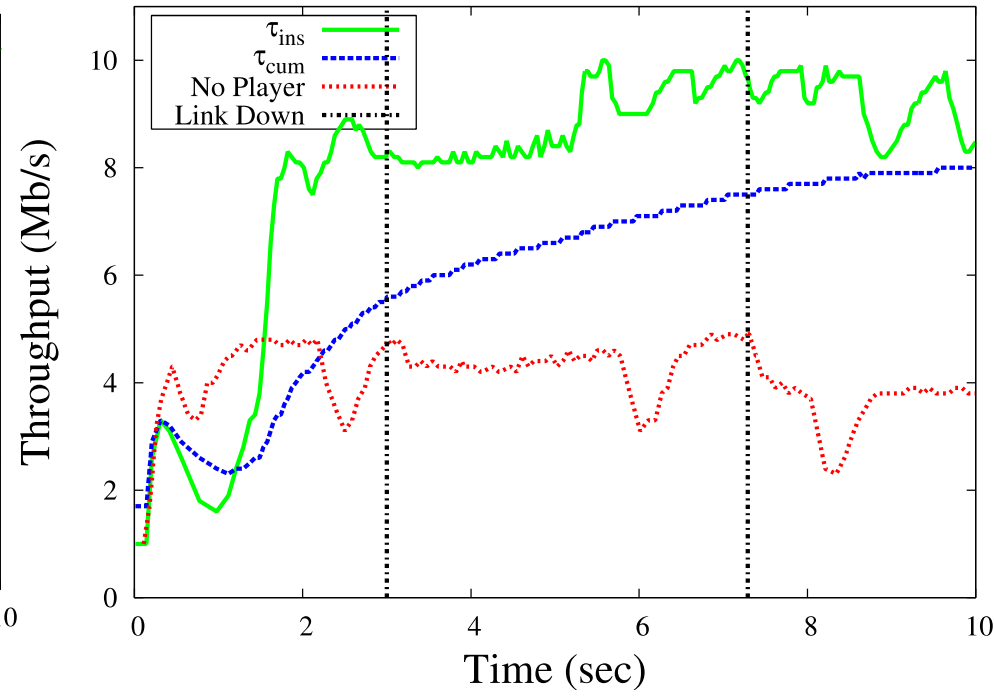

After Training
Case 7’

Best Player

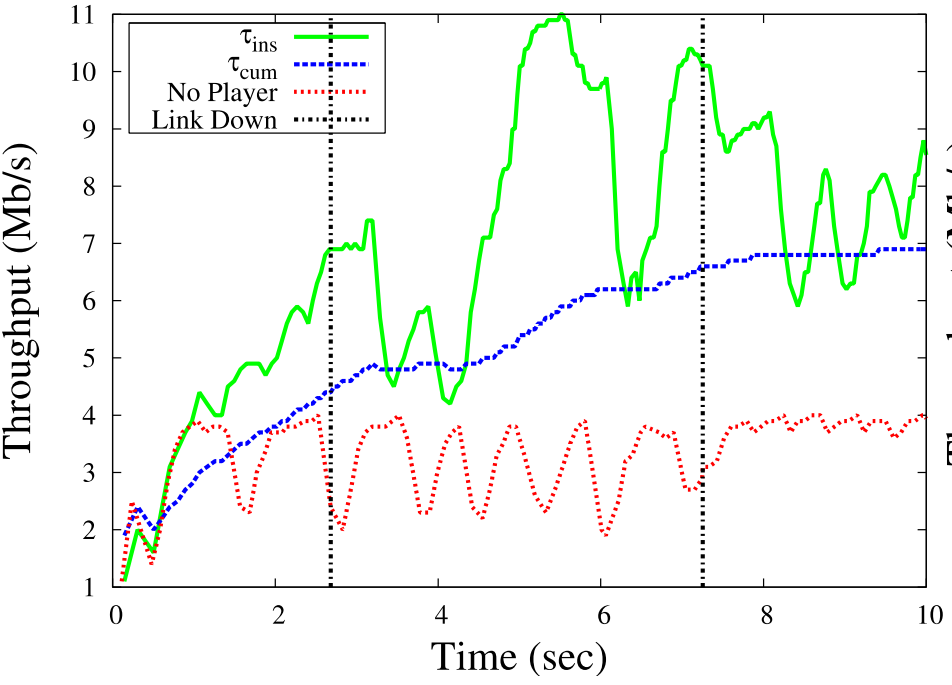

Before Training
Case 7

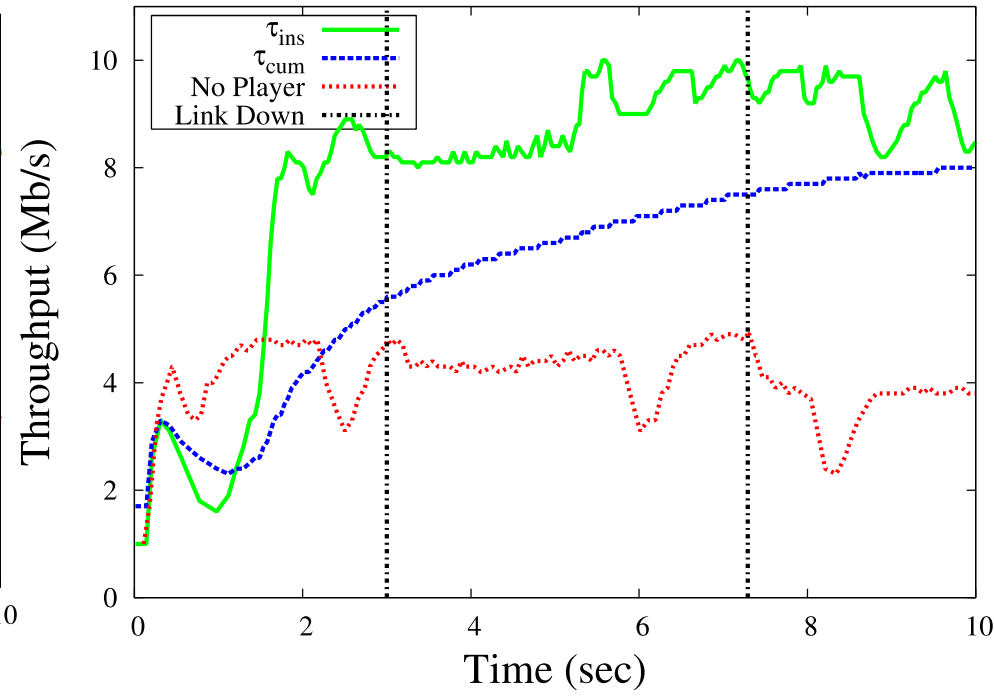

After Training
Case 7'

Worst Player

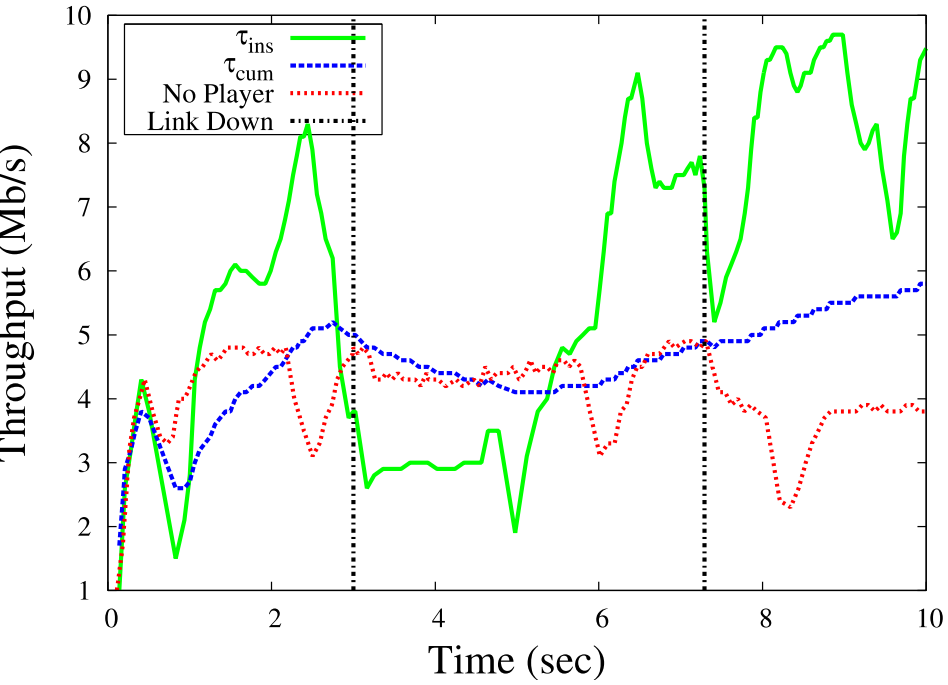

Before Training
Case 7

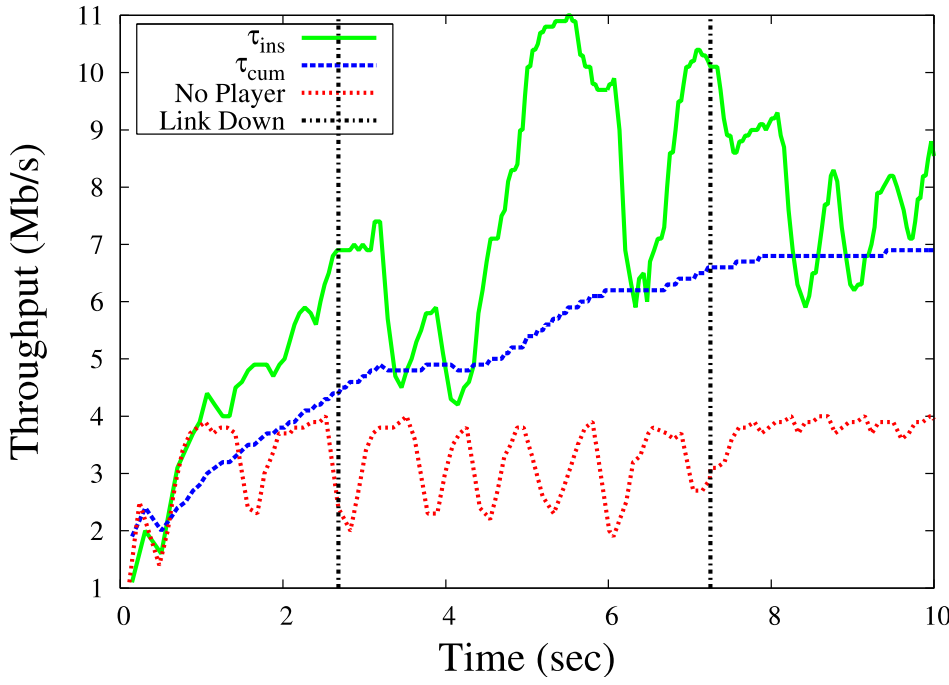

After Training
Case 7'

No Failure Case

	Before	Ratio to	After	Ratio to
No Player	*6*	*66.6*	*6*	*66.6*
Genetic Algorithms	*-*	*-*	*6.8*	*75.5*
Random Recursive	*-*	*-*	*8.5*	*94.4*
Users (Average)	*7.11*	*79*	*8.6*	*95.5*
Optimal	*9*	*100*	*9*	*100*

P-test value :0.0002

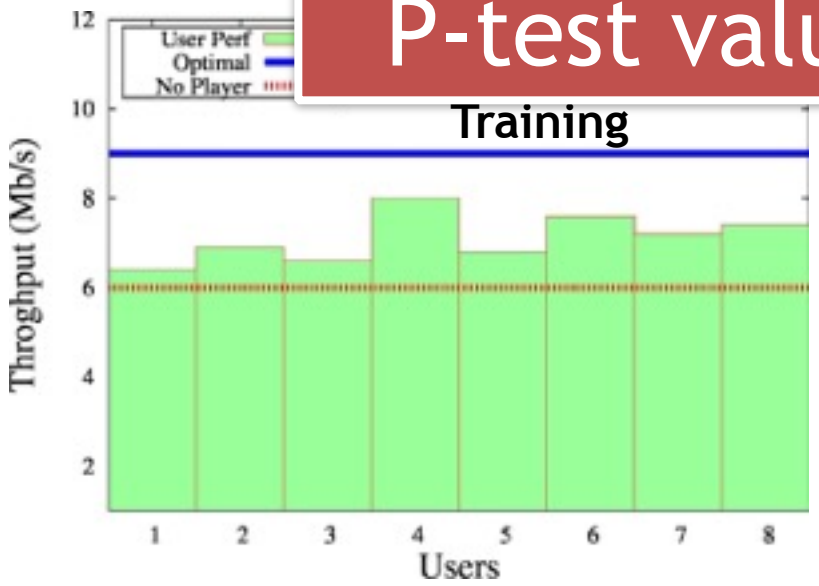

Training

After Training

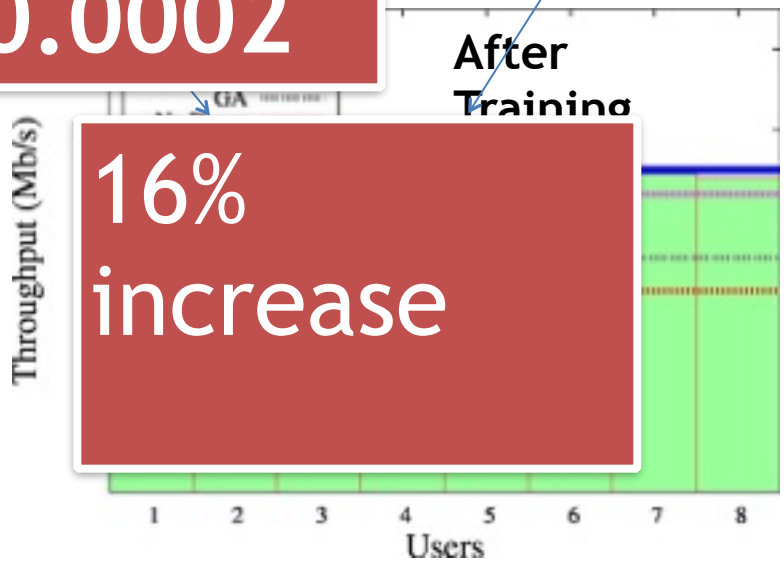

16% increase

Failure Case

	Before	Ratio to	After	Ratio to
No Player	*4*	*30.7*	*5*	*38*
Genetic Algorithms	-	-	*7.9*	*60.7*
Random Recursive			*8*	*61.5*
Users (Average			*.01*	*77*
Optimal			*3*	*100*

Users outperform

P-test value: **0.27**

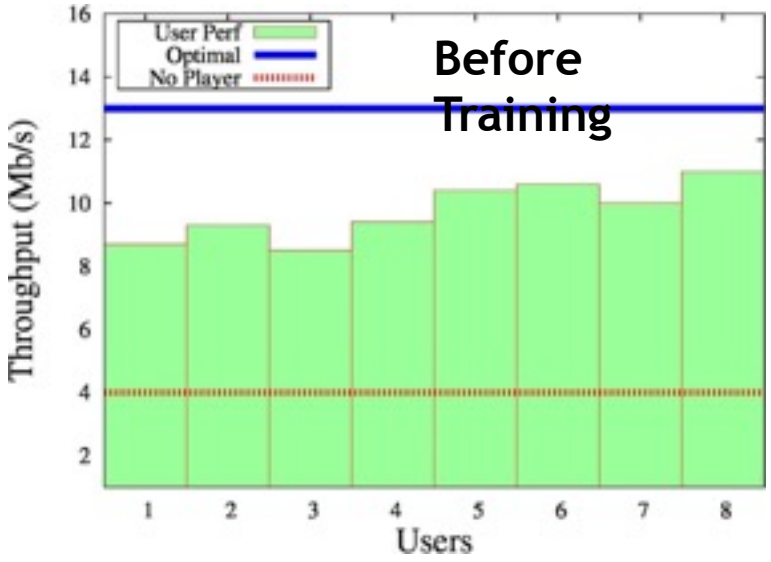

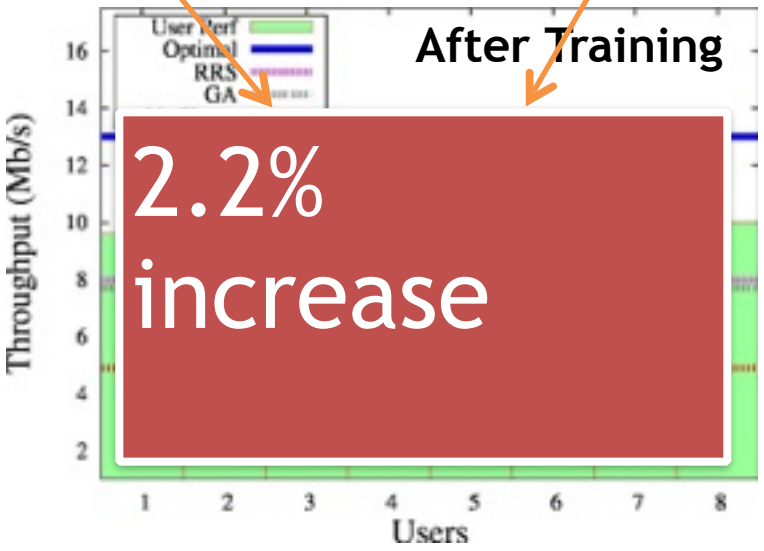

IEEE LANMAN 2011, ACM CCR 2013

Training with Mastery

- Two skills are targeted
  - ❖High bandwidth path selection
  - ❖Decoupling of flows
- 7 training scenarios→ 7 levels
- Success level is obligated to advance next level
- 5 users attended
- Took 2-3 hours on average per user

Training with Mastery

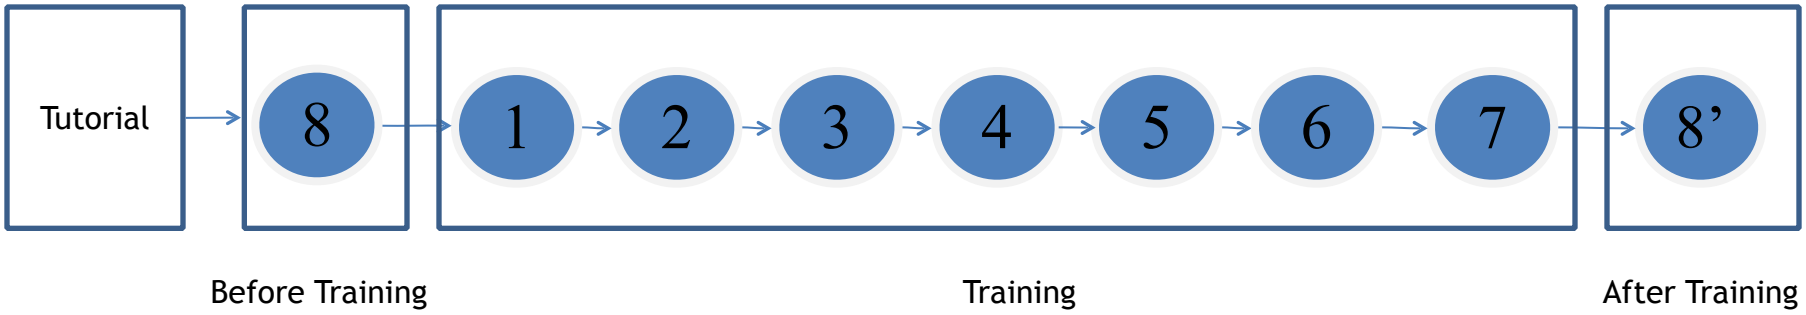

Results of Training with Mastery

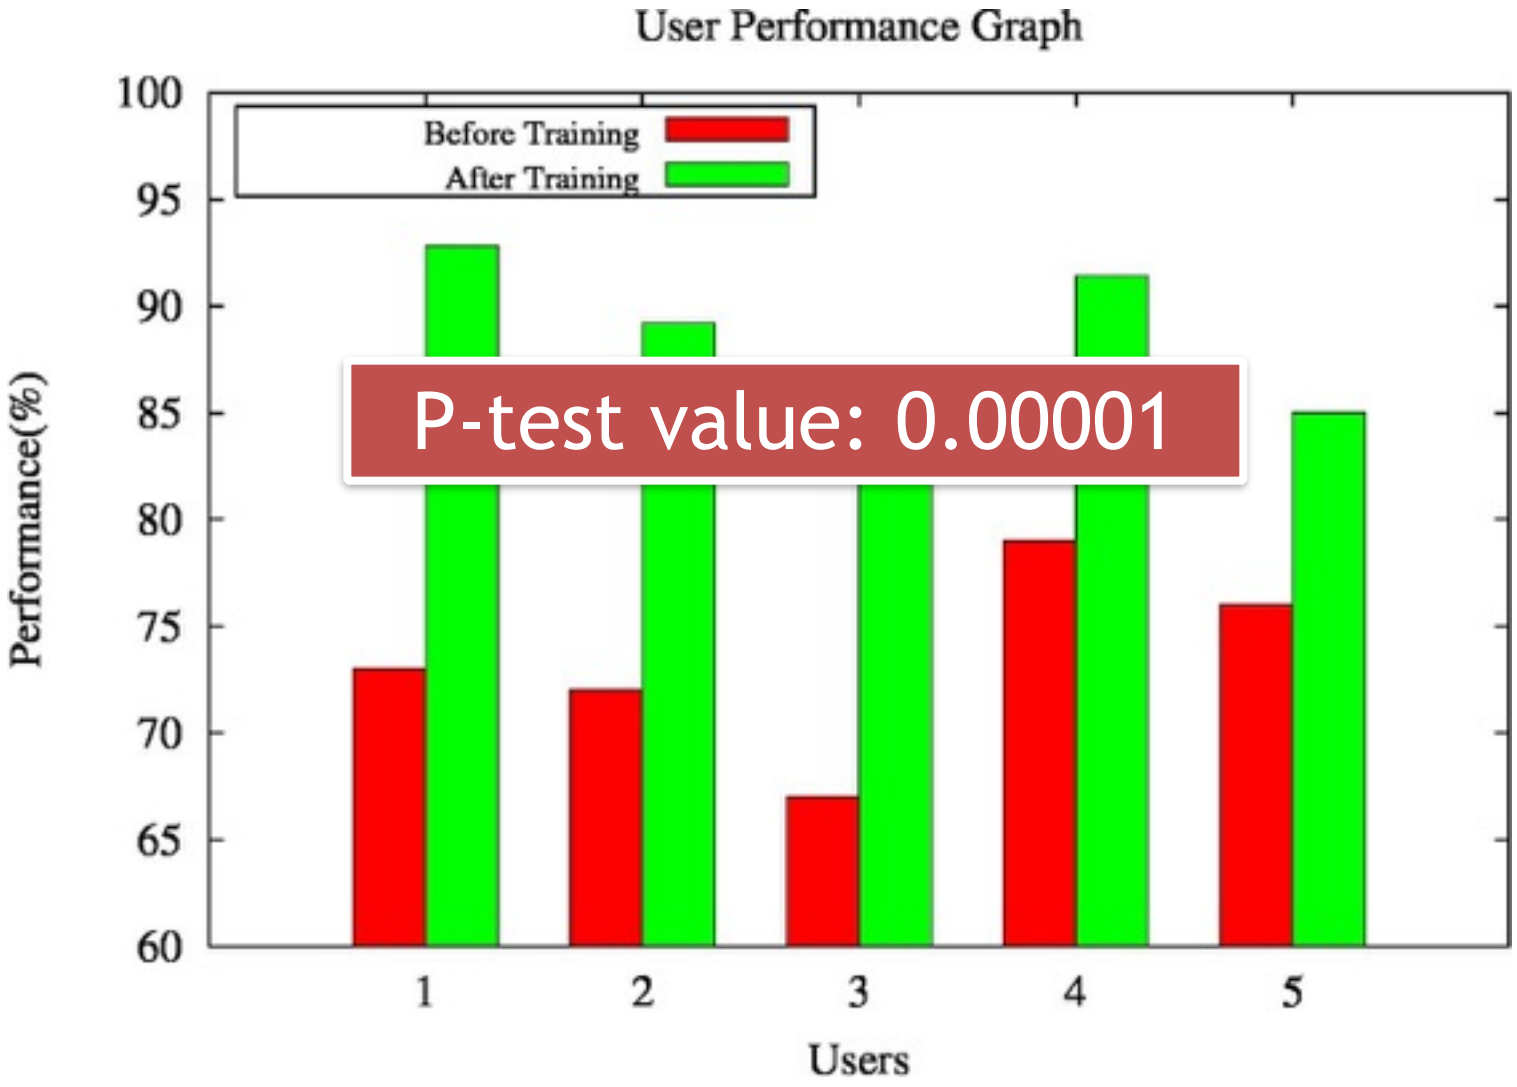

Summary

- Performance of a person in network management can be improved via our tool
  - ❖ 16% improvement → first user experiment
  - ❖ 13%- 21% improvement →second user experiment
- People outperform heuristic algorithms in case of dynamism in network
- Targeting skills and designing specific scenarios for skills lead better training
  - ❖ Success level of second user training

Future Work

- Extend quantity and quality of test cases
  - different metrics such as delay or loss
  - add what-if scenarios
  - multiple link failures
- Extend for large scale networks
- Try problems other than IGP link weight setting

- Longer term:
  - Multi-ISP games: peering wars
  - Simplified version on smartphones or web
  - Better visualization

THE END

Project Website: http://www.cse.unr.edu/~yuksem/omega.htm

Google: “omega networks unr”

AN ODD OR INEVITABLE PAIRING?

Resource management crosses disciplines

WATER, NOT LAND

IP address allocation and management as a common pool resource good

Presented by Dr. Julie Burlingame Percival, PhD
Public Policy and Political Economy

OVERVIEW OF TALK

- *Who benefits from this talk?*
- *Description of terms*
- *The fundamental dilemmas of public goods and common pool goods*
- *Possible resolutions for these dilemmas*
- *Future avenues for research*

WHO BENEFITS?

TWO AUDIENCES

Architect

Main task

Design IPv6

Main goal

Wide adoption and use of IPv6

Main frustration

Slow adoption of IPv6

Implementer

Main task

Make systems function

Main goal

Make systems function

Main frustration

IPv6 is going to break all the things

TERMS AND CONCEPTS

Types of Resources or Goods

Private

Common Pool

Toll

Public

	Exclusive	Non-Exclusive
Rivalrous	*Private (chocolate bar)*	*Common (fresh water)*
Non-rivalrous	*Club (toll road)*	*Public (state park)*

Exclusive goods : owners can prevent or allow its use.

Rivalrous goods : once used, cannot be used by someone else.

*Laws and changes in the availability of a good can shift it into being a different type

TYPES OF NON-EXCLUSIVE RESOURCES

Common Pool (CPR) goods

- *Fresh water*
- *Fish*
- *Wood*
- *Game*

Limited, but with active management is sustainable / replaceable

Public goods

- *Air*
- *Water*
- *Public parks*
- *Public infrastructure*

Durable or plentiful to the point where relatively little active management is necessary to maintain it

WHAT KIND OF GOOD IS AN IP ADDRESS?

Exclusive or non-exclusive?

- *Unique, but essentially unlimited*
- *Having an IP address is a precondition for Internet access*

Rivalrous or non-rivalrous?

It depends on the perspective

- *Non-rivalrous to end users*
- *Rivalrous to service providers*

PROPERTIES OF IP ADDRESSES

- *Used in common by a limited number of interested parties*
- *Very large, but not limitless*
- *System-wide performance affected by over-dispersion of IP addresses*

From the Implementer's perspective, IP addresses are effectively COMMON POOL RESOURCE GOODS

Architects view IP addresses as PUBLIC GOODS

FUNDAMENTAL DILEMMAS OF PUBLIC AND COMMON POOL GOODS

PUBLIC RESOURCE PROBLEMS:

Freeriders!

COMMON POOL RESOURCE PROBLEMS

- *Unequal use patterns*
- *Actors penalized for not using resource*
- *Careful use of good by individual actors not rewarded*
- *Easy to "cheat"*
- *Eventual depletion of good, also known as the Tragedy of the Commons*

THE TRAGEDY OF THE COMMONS

Actor 1 / Actor 2	Use resource	Do not use resource
Use resource	*B+2c*, *B+2c*	*B+c*, *c*
Do not use	*c*, *B+c*	*0,0*

RED represents Actor 1's strategy
BLUE represent Actor 2's strategy

Where
B= Benefit from good, c = collective cost, and B+c>B+2c>0>c

The Nash equilibrium indicates actors will act to maximize individual payoffs regardless of the consequence to that resource
This game is a variant of "The Prisoner's Dilemma"

WHO CAN PREVENT THE TRAGEDY OF THE COMMONS?

- *Government (public)*
- *The Firm (private)*
- *Self-governed*

PUBLIC VS PRIVATE CONTROL OF THE COMMONS

Government (public)

- *Resource allocation determined by non-user*
- *Slow to respond to new conditions*
- *Overly cautious. High penalties for exploitive use*
- *Inefficient*

The Firm (private)

- *Resource allocation determined by single user*
- *Quick to respond to new conditions*
- *Overly aggressive*
- *Efficient until resource is depleted*

SELF-GOVERNED CONTROL OF THE COMMONS

- *Interested only in the maintenance of a renewable resource*
- *Multi-stakeholders band together to agree on rules for usage*
- *Most responsive to asset management for parties utilizing unequal amounts of the resource*
- *Elinor Ostrom won a Nobel in Economics in part for demonstrating how self-governed cooperation of actors can shift the strategic game outcomes for a CPR good in Governing The Commons*

RESOLVING IP ALLOCATION DILEMMAS

TRAITS ASSOCIATED WITH "STRONG" CPR SELF-GOVERNANCE

- *Clear boundaries*
- *Congruence between rules and local conditions*
- *Collective-choice arrangements*
- *Effective monitoring*
- *Graduated sanctions*
- *Conflict-resolution arrangements*
- *Some recognized rights to organize or bring petitions*
- *Nested system arrangements*

SOME DILEMMAS IN IP ADDRESS ALLOCATION

1. *Not enough IP addresses*
2. *Workarounds violate core architectural protocol*
3. *Router memory and speed limits growth and overall functionality*
4. *“Hoarding” and underuse of available addresses*
5. *“Fair” methods of distribution*

TAKEAWAY

While the "IP addresses are like home addresses" is useful for describing the concept of IP addresses to end-users, the properties and usage of IP addresses more closely model that of community resources like clean water.

Experimental research can model and test whether these hypotheses about how people use IP addresses in theory and in practice are reasonable.

FUTURE RESEARCH

FUTURE RESEARCH

1) *Create models that simulate and test the theory: agent based modeling can potentially test different usage strategies and model actor behavior*
2) *Compare simulation models to historical routing table data or current routing table data*
3) *Experiment! Test and compare route dispersion from ABM model predictions in experimental and control groups under a variety of different conditions*

Pacific Research Platform

June 1, 2015

NANOG64
San Francisco

Corporation for Education Network Initiatives in California (CENIC):

To advance education and research throughout California by providing the world-class network essential for innovation, collaboration and economic growth – connecting California to the world.

CENIC

Connecting California

K-12 High Speed Network

California's Research & Education Network

CENIC is a 501(c)3 created to serve California's K-20 research & education institutions with cost-effective, high-bandwidth networking

Five Charter Associates: California Community Colleges, California K-12 System, California State University System, Private Universities, and the University of California System

http://www.cenic.org

CENIC: California's Research & Education Network

- **3,800+ miles of optical fiber**
- Members in all 58 counties connect via fiber-optic cable or leased circuits from telecom carriers.
- **Nearly *10,000* sites connect to CENIC**
- ***20,000,000* Californians use CENIC each day**
- Governed by members on the segmental level

CENIC: California's Research & Education Network

Three networks operate simultaneously as independent layers on a single infrastructure.

- **CalREN-Digital California (DC) / AS2152:** commodity Internet access, e-mail, web browsing, videoconferencing, etc.
- **CalREN-High-Performance Research (HPR) / AS2153:** high-performance research for big-science, inter-institutional collaborations
- **CalREN-eXperimental Developmental (XD):** research on the network itself

Pacific Wave: CENIC & PNWGP

- A joint project between CENIC and the Pacific Northwest Gigapop (PNWGP)
- Operated in collaboration with the University of Southern California and the the University of Washington

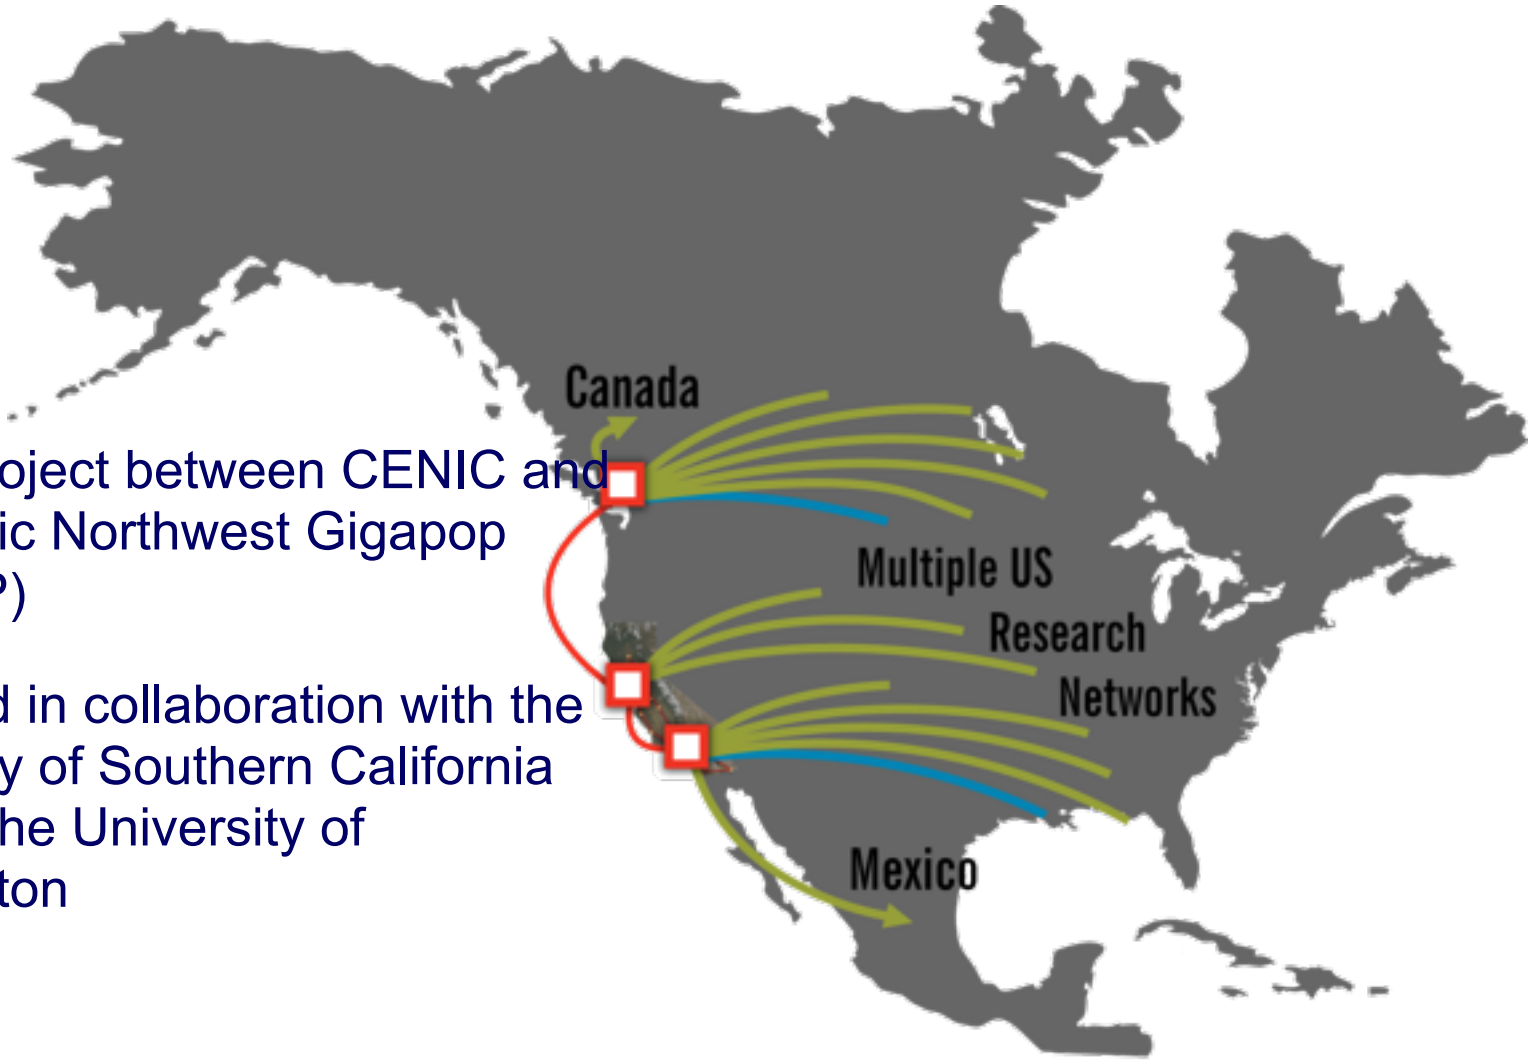

http://www.pacificwave.net/

Pacific Wave: enables *worldwide* collaboration

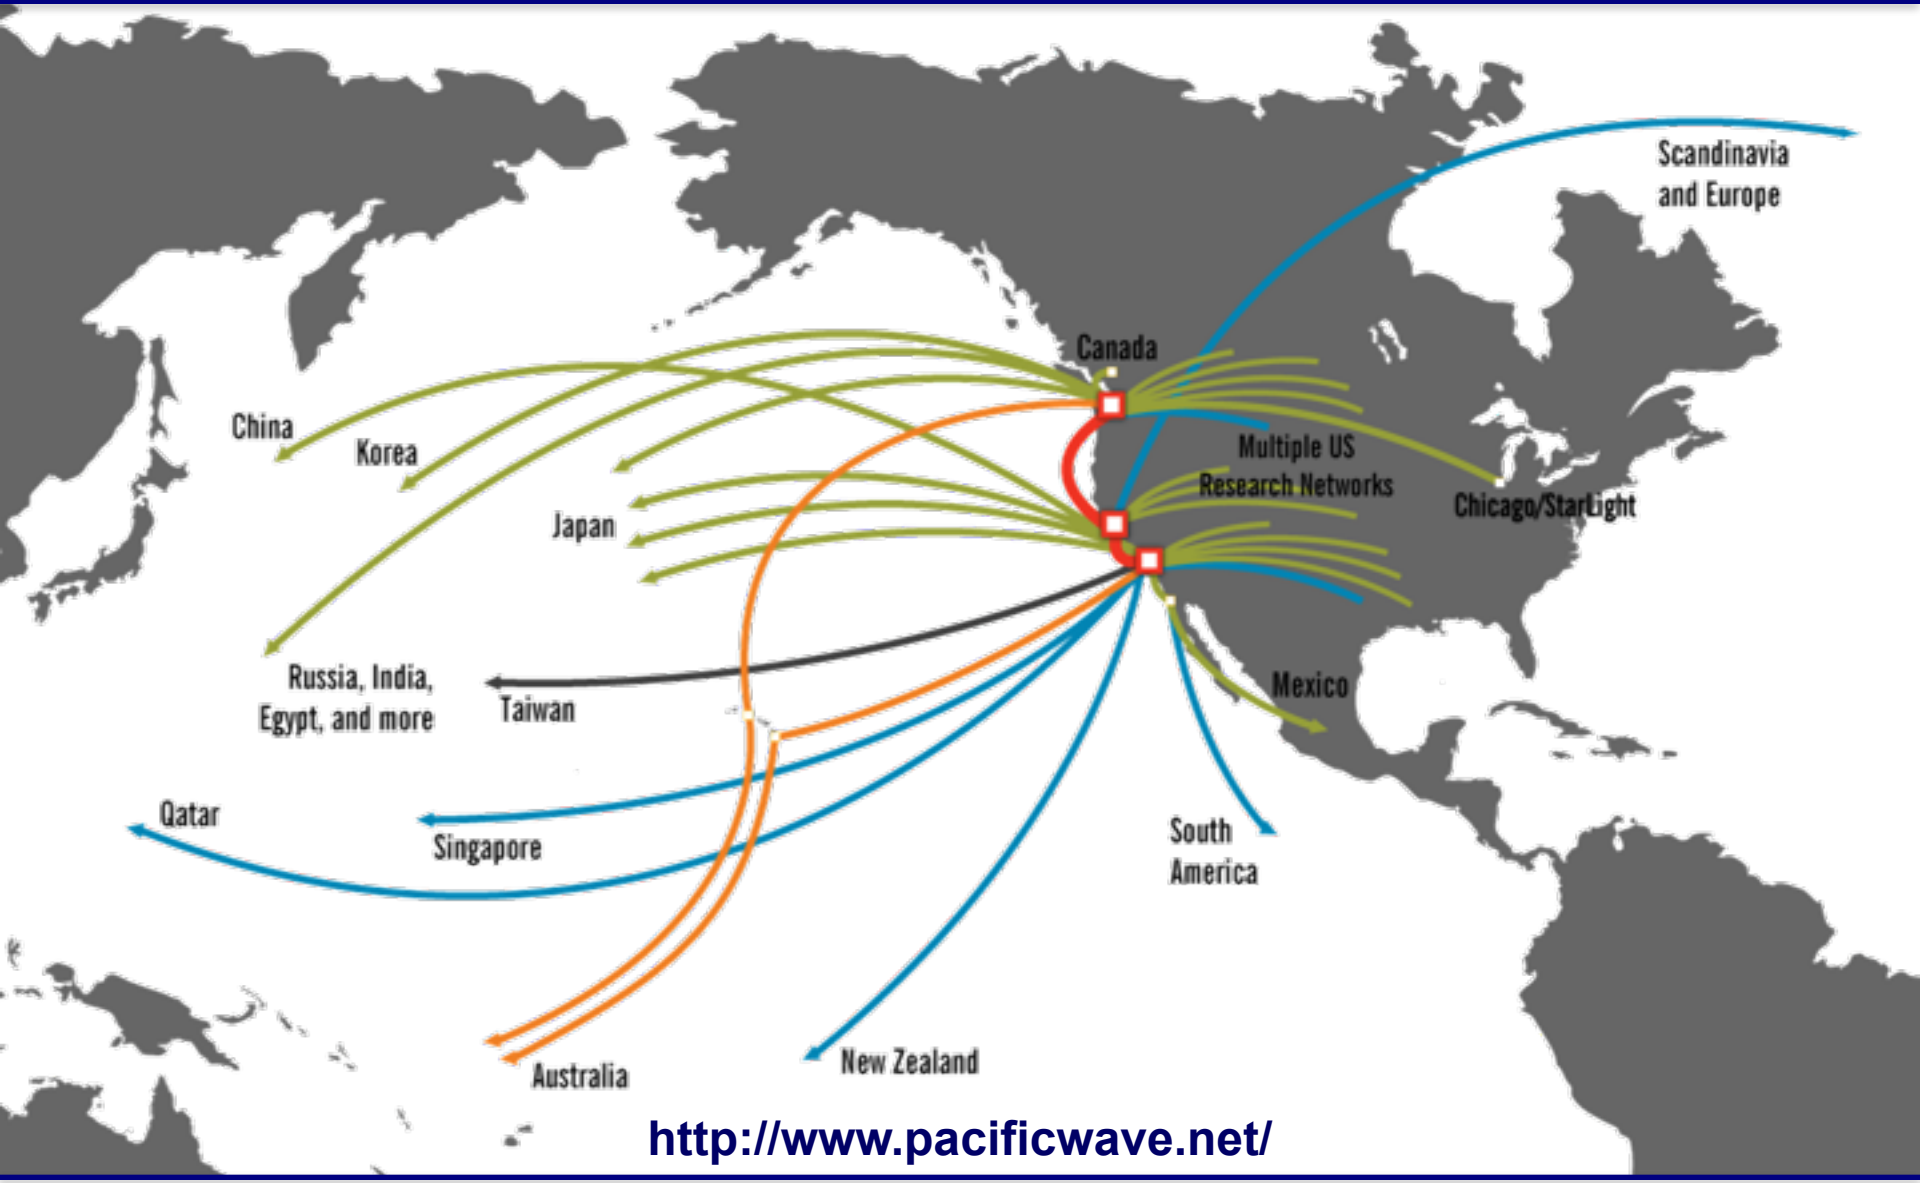

2014-2015 CENIC INITIATIVES

- California Community Colleges
- K12 Last Mile
- California Public Libraries
- 100G Backbone
- Pacific Research Platform

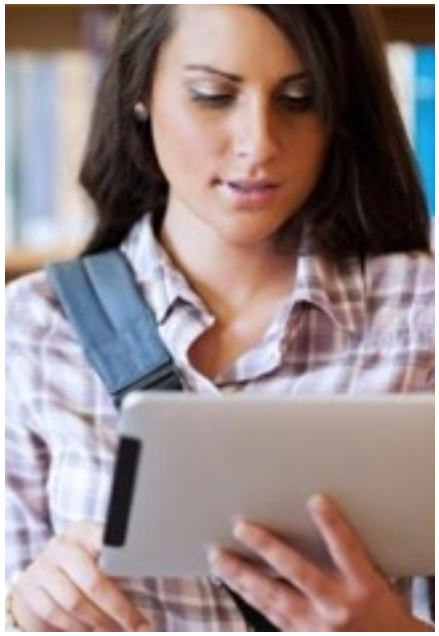

CENTER FOR IT RESEARCH

CENIC

CITRIS
citris-uc.org

it²

CITRIS and Calit2…

Center for Information Technology Research in the Interest of Society

&

California Institute for Telecommunications & Information Technology

- Governor Gray Davis Institutes of Science and Innovation since 2001
- Multi-campus, multidisciplinary research institutes
- Charged with creating IT solutions for society's most pressing challenges

CITRIS and Calit2…

- Together we cover 6 of 10 UC campuses
- Major research initiatives in
  - Health
  - Energy and the Environment
  - Robotics
  - Connected Communities
  - Nanotechnology

Pacific Research Platform

Abstract

The Pacific Research Platform is a project to forward the work of advanced researchers and their access to technical infrastructure, with a vision of connecting all the National Science Foundation Campus Cyberinfrastructure grants (NSF CC-NIE & CC-IIE) to research universities within the region, as well as the Department of Energy (DOE) national labs and the San Diego Supercomputer Center (SDSC).

Science Drivers

Particle Physics

Astronomy and Astrophysics

Biomedical

Earth Sciences

Scalable Visualization, Virtual Reality, and Ultra-Resolution Video

Science Drivers

Particle Physics Data Analysis

- The Large Hadron Collider (LHC). Run 2 will have ~2x the energy, generating ~10x the data volume of Run 1.

Astronomy and Astrophysics Data Analysis

- Includes two data-intensive telescope surveys that are precursors to the Large Synoptic Survey Telescope (LSST)
  - *Intermediate Palomar Transient Factory (iPTF)*
  - *Dark Energy Spectroscopic Instrument (DESI)*
- Galaxy Evolution
  - *Southern California Center for Galaxy Evolution (CGE)*
  - *Assembling Galaxies of Resolved Anatomy (AGORA)*
- Gravitational Wave Astronomy
  - *The Laser Interferometer Gravitational-Wave Observatory (LIGO)*

Biomedical Data Analysis

Cancer Genomics Hub (CG Hub) and Cancer Genomics Browser
Microbiome and Integrative 'Omics
Integrative Structural Biology

Science Drivers (2)

Earth Sciences Data Analysis

- Data Analysis and Simulation for Earthquakes and Natural Disasters
  Pacific Earthquake Engineering Research Center (PEER)
- Climate Modeling
  National Center for Atmospheric Research (NCAR)
  University Corporation for Atmospheric Research (UCAR)
- California/Nevada Regional Climate Data Analysis
  California Nevada Climate Applications Program (CNAP)
- CO2 Subsurface Modeling

Scalable Visualization, Virtual Reality, and Ultra-Resolution Video

Cultural Heritage Data
Networked Scalable Visualization
Virtual Reality Systems
Ultra-Resolution Video Systems

The Science DMZ* in 1 Slide

Consists of **three key components**, all required:

- "Friction free" network path
  - Highly capable network devices (wire-speed, deep queues)
  - Virtual circuit connectivity option
  - Security policy and enforcement specific to science workflows
  - Located at or near site perimeter if possible
- Dedicated, high-performance Data Transfer Nodes (DTNs)
  - Hardware, operating system, libraries all optimized for transfer
  - Includes optimized data transfer tools such as Globus Online and GridFTP
- Performance measurement/test node
  - perfSONAR
- Engagement with end users

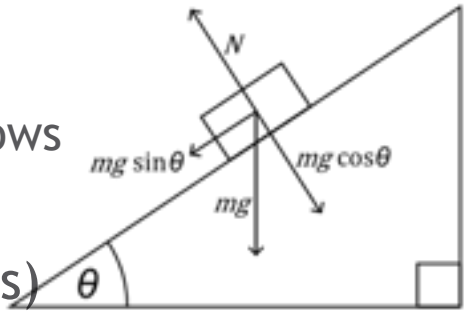

© 2013 Wikipedia

globus

perfSONAR

Details at http://fasterdata.es.net/science-dmz/

* ***Science DMZ*** is a trademark of The Energy Sciences Network (ESnet)

ESnet

Science DMZ Superfecta: Engagement

Engagement
- Partnerships
- Education & Consulting
- Resources & Knowledgebase

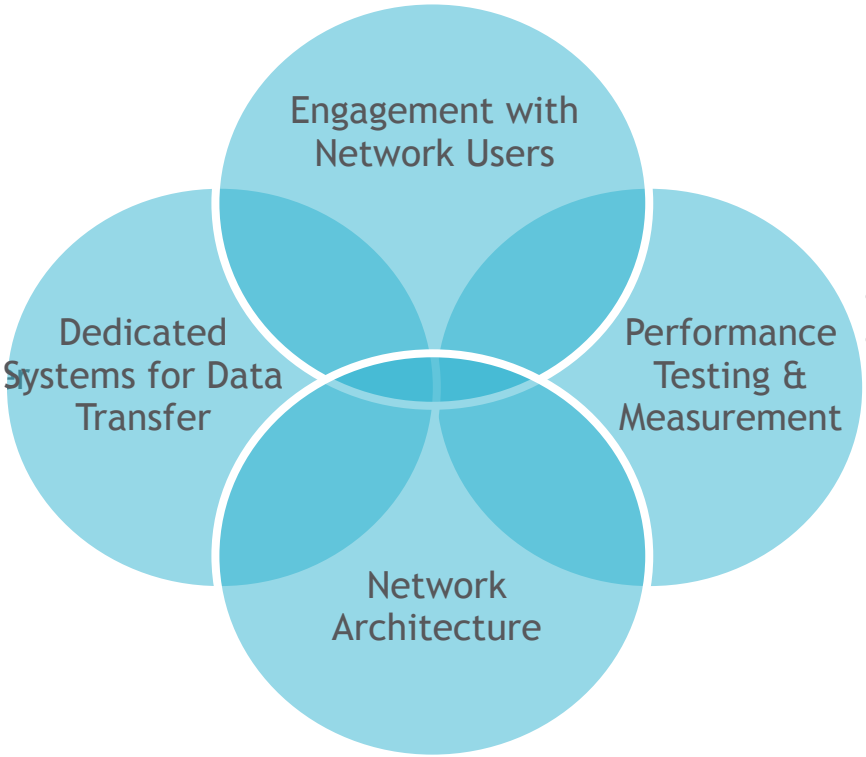

Data Transfer Node
- High performance
- Configured for data transfer
- Proper tools

perfSONAR
- Enables fault isolation
- Verify correct operation
- Widely deployed in ESnet and other networks, as well as sites and facilities

Science DMZ
- Dedicated location for DTN
- Proper security
- Easy to deploy - no need to redesign the whole network

ESnet

NSF Funding Has Enabled Science DMZs at Over 100 U.S. Campuses

- **2011 ACCI Strategic Recommendation to the NSF #3:**
  - **NSF should create a new program funding high-speed (currently 10 Gbps) connections from campuses to the nearest landing point for a national network backbone. The design of these connections must include support for dynamic network provisioning services and must be engineered to support rapid movement of large scientific data sets."**
  - **- pg. 6, NSF Advisory Committee for Cyberinfrastructure Task Force on Campus Bridging, Final Report, March 2011**
    - **www.nsf.gov/od/oci/taskforces/TaskForceReport_CampusBridging.pdf**
  - **Led to Office of Cyberinfrastructure CC-NIE RFP March 1, 2012**
- **NSF's Campus Cyberinfrastructure – Network Infrastructure & Engineering (CC-NIE) Program**
  - **>130 Grants Awarded So Far (New Solicitation Open)**
    - **Roughly $500k per Campus**

Next Logical Step-Interconnect Campus Science DMZs

Pacific Research Platform Strategic Arc

Build upon Pacific Wave as a backplane for data-intensive science

- High performance data movement provides capabilities that are otherwise unavailable to scientists
- Integrating Science DMZs across the West Coast
- This capability is extensible, both regionally and nationally

Goal: scientists can get the data they need, where they need it, when they need it

- PRPv0: a proof of concept experiment to develop and inform requirements for future work.
- Engage with scientists to map their research on to the Pacific Research Platform

PRPv0 -- An experiment including:

Caltech
CENIC / Pacific Wave
ESnet / LBNL
NASA Ames / NREN
San Diego State University
SDSC
Stanford University
University of Washington
USC

UC Berkeley
UC Davis
UC Irvine
UC Los Angeles
UC Riverside
UC San Diego
UC Santa Cruz

What have we done

PRPv0 concentrated on the regional aspects of the problem. There are lots of parts to the research data movement challenge. This experiment mostly looked at the inter-campus piece . Over a 10-week period, lots of network and HPC staff at lots of sites collaborated to

- Build a mesh of perfSONAR instances to instrument the network
- Implement MaDDash -- Measurement and Debugging Dashboard
- Deploy Data Transfer Nodes (DTNs)
- Perform GridFTP file transfers to quantify throughput
- Activate an ad-hoc, partial BGP peering mesh across a fabric of 100G links to demonstrate the potential of networks with burst capacity greater than that of a single DTN
- Identify some specific optimizations needed
- Fix a few problems in pursuit of gathering illustrative data
- Identify anomalies for further investigation

MaDDash of perfSONAR throughput and loss

- Performance for nodes that are close is better than for nodes that are far away
- Network problems that manifest over a distance may not manifest locally

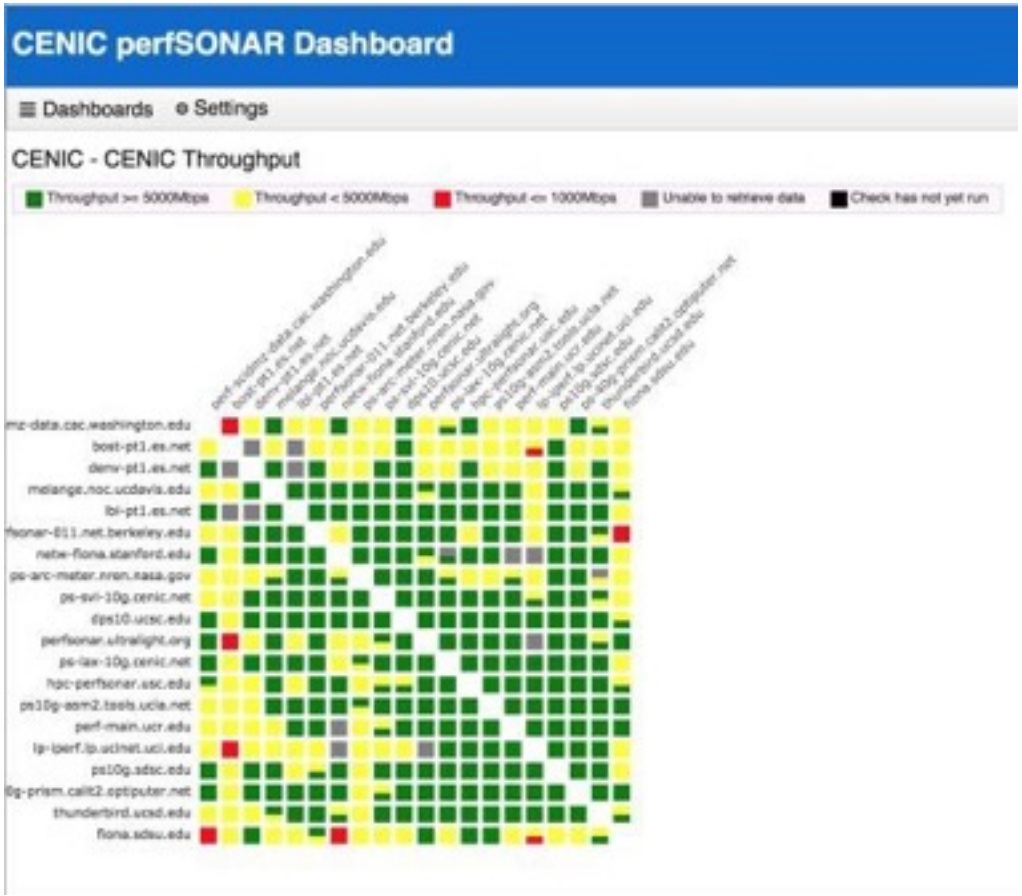

Science DMZ Data Transfer Nodes Can Be Inexpensive PCs Optimized for Big Data

- **FIONA – Flash I/O Node Appliance**
  - Combination of Desktop and Server Building Blocks
  - US$5K - US$7K
  - Desktop Flash up to 16TB
  - RAID Drives up to 48TB
  - 10GbE/40GbE Adapter
  - Tested speed 40Gbs
  - Developed Under UCSD CC-NIE Prism Award by UCSD's
    - Phil Papadopoulos
    - Tom DeFanti
    - Joe Keefe

SDSC

For More on Science DMZ DTNs See:
https://fasterdata.es.net/science-dmz/DTN/

MaDDash of GridFTP transfers

- DTNs loaded with Globus Connect Server suite to obtain GridFTP tools.
- cron-scheduled transfers using globus-url-copy.
- ESnet-contributed script parses GridFTP transfer log and loads results in an esmond measurement archive.

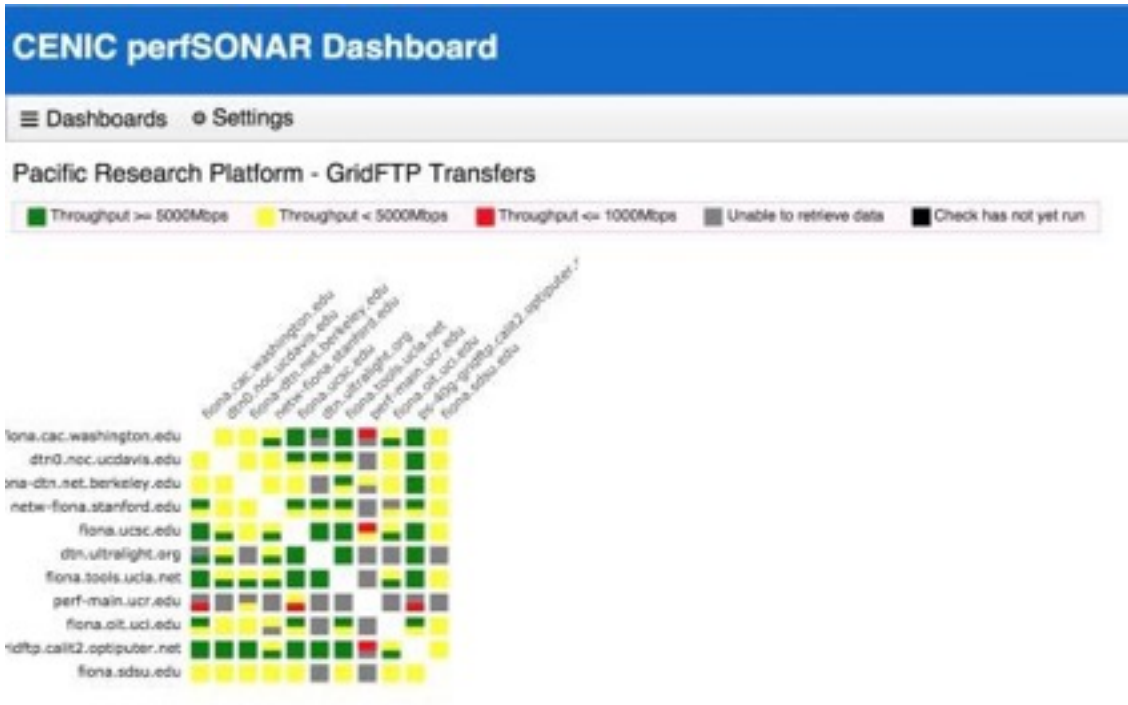

Next Steps and Near-term Goals

- Migrate from experiment to persistent infrastructure as part CalREN HPR 100G Layer 3 upgrade
- Expand perfSONAR measurement and test infrastructure
- Apply 2015 funding from UC Office of the President toward a DTN deployment to include all 10 UC campuses and compliment existing DTN deployments at Caltech, Stanford, USC, and University of Washington
- Incorporate federated authentication for access to resources
- Engage with scientists to begin to map their research collaborations across the Pacific Research Platform
- Work with campus IT organizations to make “last mile” connections between researchers and the Science DMZ

Longer-term Goals

- An Integrated West Coast Science DMZ for Data-Intensive Research
- Advocate for similar projects based on ESnet's ScienceDMZ model
- Science DMZ interoperability / integration across regions, nationally, and internationally
- SDN/SDX, …
- Commercial services – Amazon AWS, Microsoft Azure, Google, …

Pacific Research Platform: A Regional Science DMZ

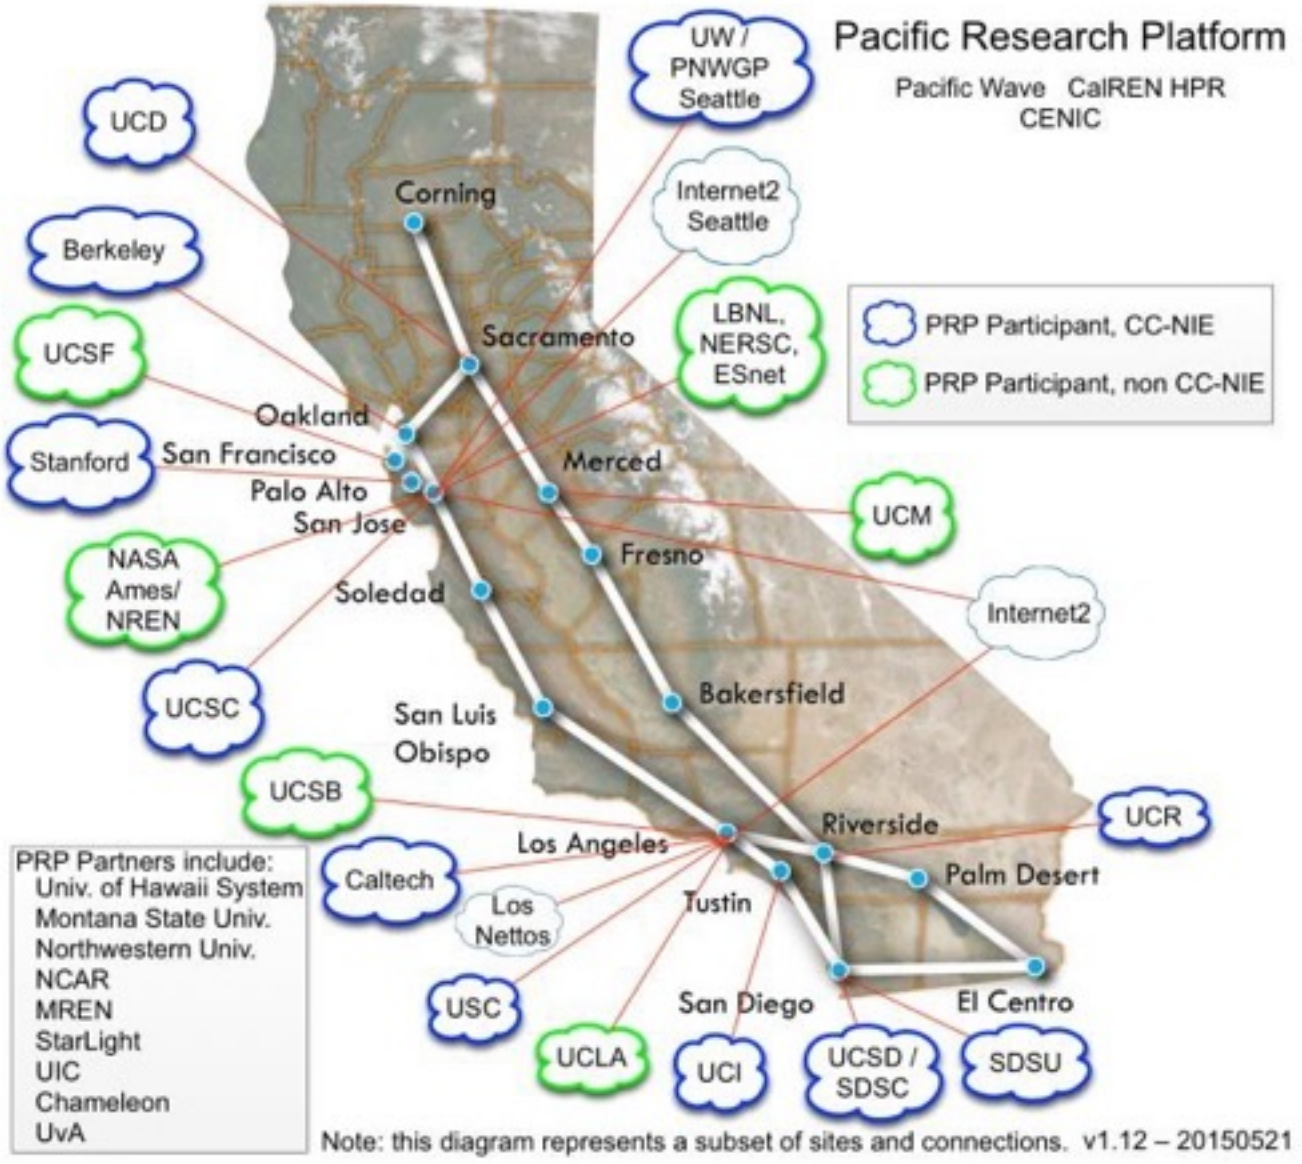

Links

- ESnet fasterdata knowledge base
  - http://fasterdata.es.net/
- Science DMZ paper
  - http://www.es.net/assets/pubs_presos/sc13sciDMZ-final.pdf
- Science DMZ email list
  - https://gab.es.net/mailman/listinfo/sciencedmz
- perfSONAR
  - http://fasterdata.es.net/performance-testing/perfsonar/
  - http://www.perfsonar.net

Pacific Research Platform

Questions?

Research and Education Track

NANOG 64

San Francisco

1 June 2015

Introduction to the R&E Track

- Borne out of Research Forum
  - Ongoing research, solicit avenues for further research.
- Internet2 Joint Techs
  - Now Internet2 “Tech Exchange”
  - More formal, only once per year

Upcoming Meetings

- Quilt – Member Meeting – Austin, 28 September – 1 October 2015
- Internet2 – Tech Exchange – Cleveland (Case Western Reserve) 4-7 October 2015
- NANOG 65 – Montreal – 5-7 October 2015
- Anyone know what's up with NetGurus?

Upcoming Meetings

- Quilt – Member Meeting – Austin, 28 September – 1 October 2015
- Internet2 – Tech Exchange – Cleveland (Case Western Reserve) 4-7 October 2015
- NANOG 65 – Montreal – 5-7 October 2015

OOPS!!!

Today's Agenda

- Murat Yuksel: Training Network Administrators in a Game-Like Environment
- Julie Percival: Water, Not Land
- Michael Smitasin: Evaluating Network Buffer Size requirements for Very Large Data Transfers
- John Hess and Camille Crittenden: The Pacific Research Platform
- Michael Sinatra: Science DMZ Security

ESnet
ENERGY SCIENCES NETWORK

Science DMZ as a Security Architecture

Nick Buraglio
Michael Sinatra
Network Engineers, ESnet
Lawrence Berkeley National Laboratory

CENIC 2015
Irvine, CA
March 11, 2015

U.S. DEPARTMENT OF ENERGY
Office of Science

BERKELEY LAB

Motivations

- You have a Science DMZ
- You need a Science DMZ
- Adding visibility is essential for accountability
- Timely mitigation of issues is required
- Automated mitigation is highly desirable
- Close to real time responses
- Providing confidentiality, accountability and integrity on an open perimeter network is the exception and not the rule
- You have research systems that are hard to manage and/or hard to secure

The Science DMZ* in 1 Slide

Consists of **three key components**, all required:

- "Friction free" network path
  - Highly capable network devices (wire-speed, deep queues)
  - Virtual circuit connectivity option
  - Security policy and enforcement specific to science workflows
  - Located at or near site perimeter if possible
- Dedicated, high-performance Data Transfer Nodes (DTNs)
  - Hardware, operating system, libraries all optimized for transfer
  - Includes optimized data transfer tools such as Globus Online and GridFTP
- Performance measurement/test node
  - perfSONAR
- Engagement with end users

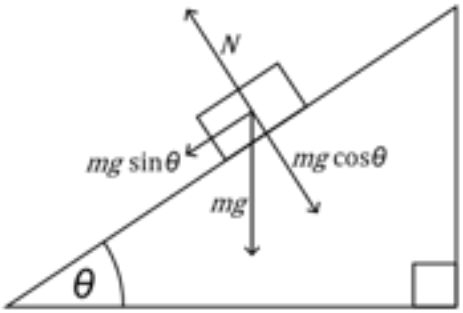

globus

perfSONAR

Details at http://fasterdata.es.net/science-dmz/

* ***Science DMZ*** is a trademark of The Energy Sciences Network (ESnet)

Myths about the Science DMZ

- “ESnet invented the Science DMZ.”

Myths about the Science DMZ

- **"ESnet invented the Science DMZ."**
- **FALSE:** You invented the Science DMZ. ESnet examined the practices that were already in place and evolving within the community and generalized them into a replicable model.

Myths about the Science DMZ

- **"ESnet invented the Science DMZ."**
- **FALSE:** You invented the Science DMZ. ESnet examined the practices that were already in place and evolving within the community and generalized them into a replicable model.

- **"The purpose of a Science DMZ is to get around or avoid firewalls and other security controls."**

Myths about the Science DMZ

- **"ESnet invented the Science DMZ."**
- **FALSE:** You invented the Science DMZ. ESnet examined the practices that were already in place and evolving within the community and generalized them into a replicable model.

- **"The purpose of a Science DMZ is to get around or avoid firewalls and other security controls."**
- **VERY FALSE:** The purpose of the Science DMZ is to match controls (security and otherwise) with the actual thing that's being protected, while maximizing the particular functionality of the network (in this case, data transfer). As such, the Science DMZ is a security architecture: The separation afforded by the Science DMZ allows finer-grained security controls tailored to the specific risks present in different parts of the network.

Myths about the Science DMZ

- “ESnet invented high-performance networking.”

Myths about the Science DMZ

- **"ESnet invented high-performance networking."**
- **FALSE:** Smart people in our community have been doing high-performance networking for years. ESnet examined the practices that were effective, and generalized them into a replicable model. That model is the Science DMZ. ESnet developed the *concept* of the Science DMZ, but many of the practices came from the community.

Myths about the Science DMZ

- **“ESnet invented high-performance networking.”**
- **FALSE:** Smart people in our community have been doing high-performance networking for years. ESnet examined the practices that were effective, and generalized them into a replicable model. That model is the Science DMZ.

- **“The purpose of a Science DMZ is to get around or avoid firewalls and other security controls.”**

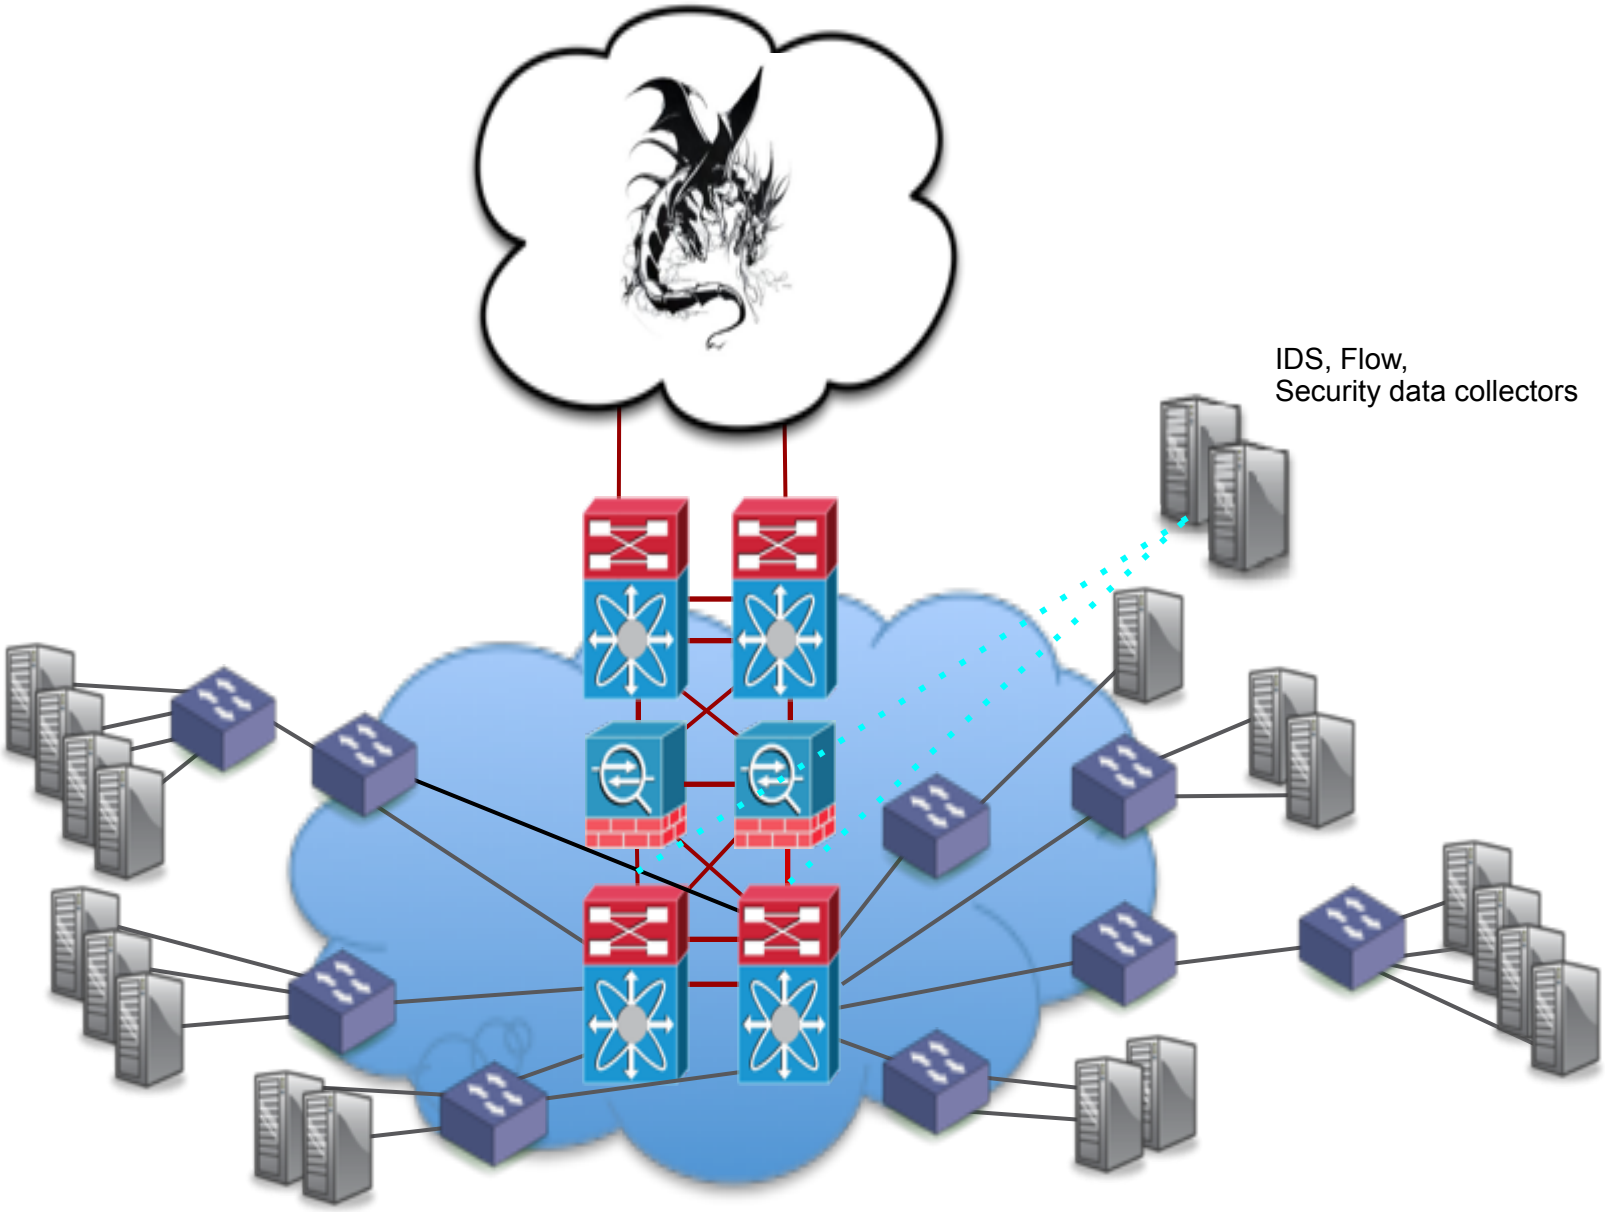
IDS, Flow,
Security data collectors

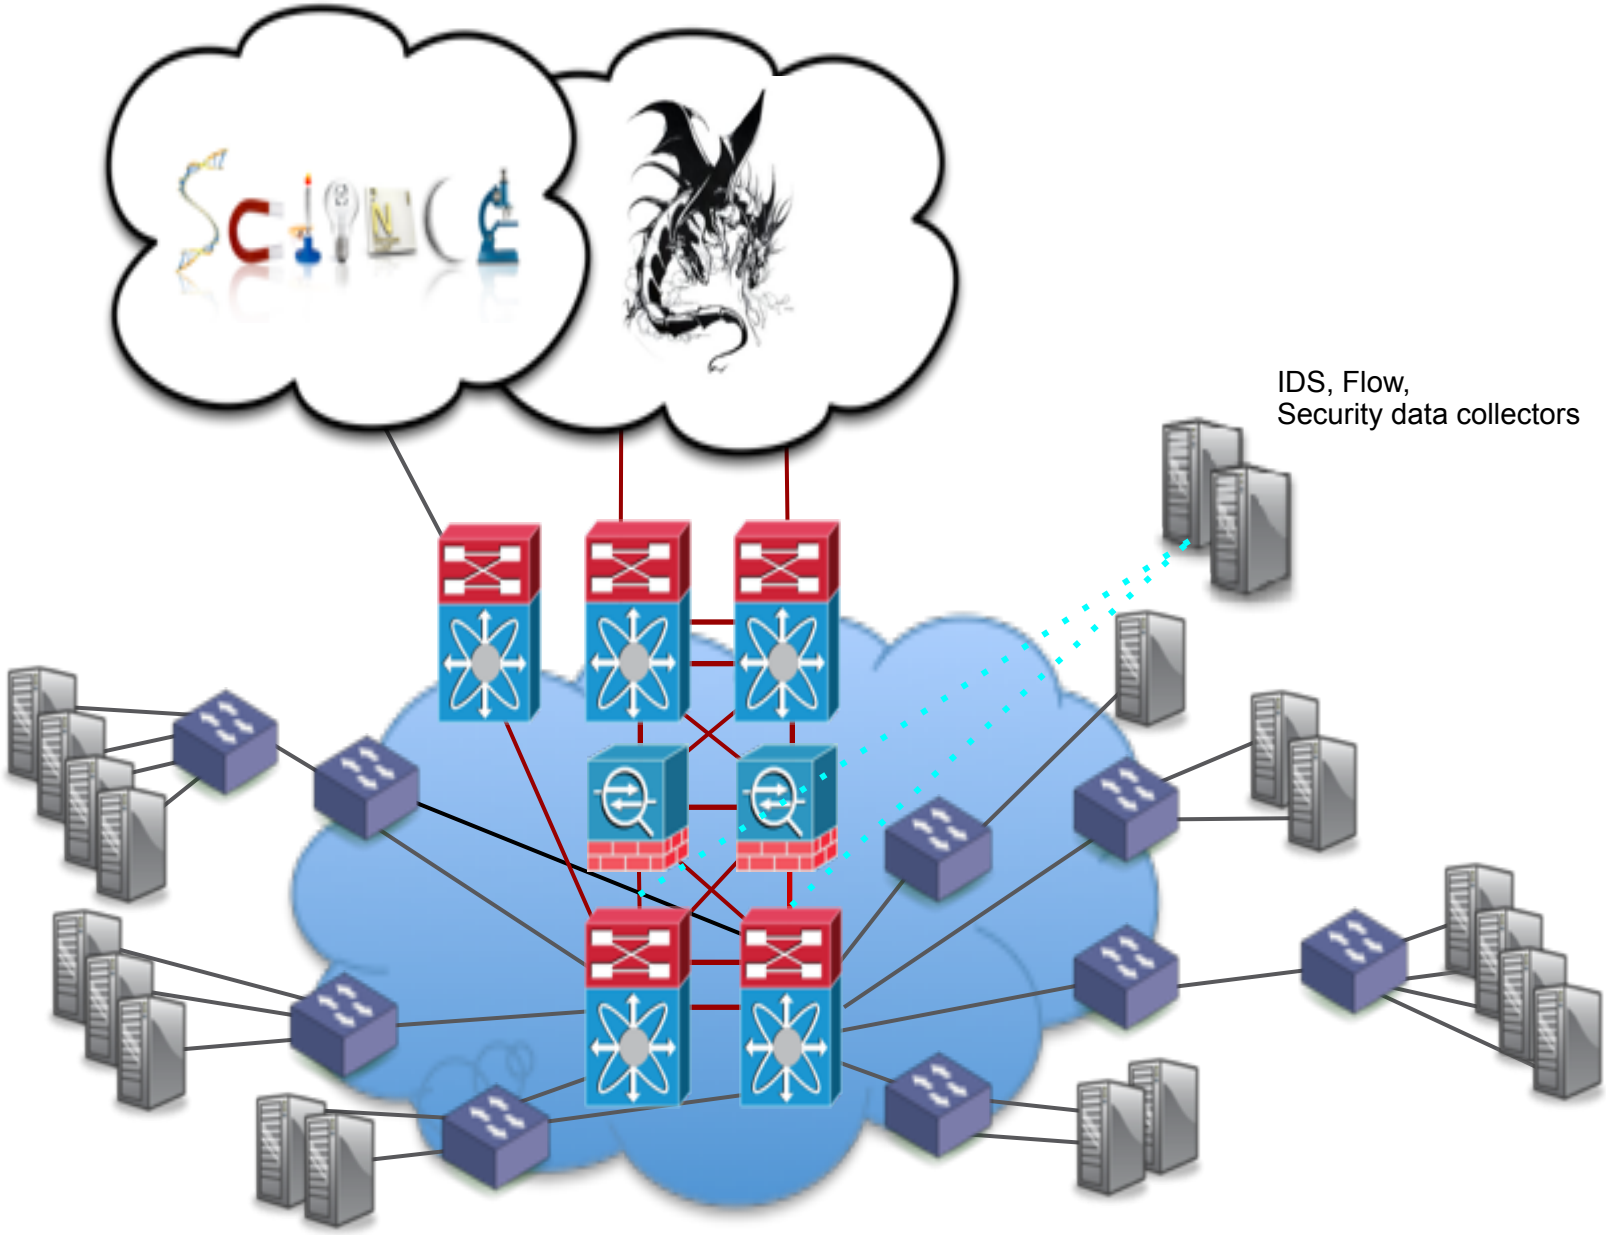

Science Image from http://www.science.fau.edu/

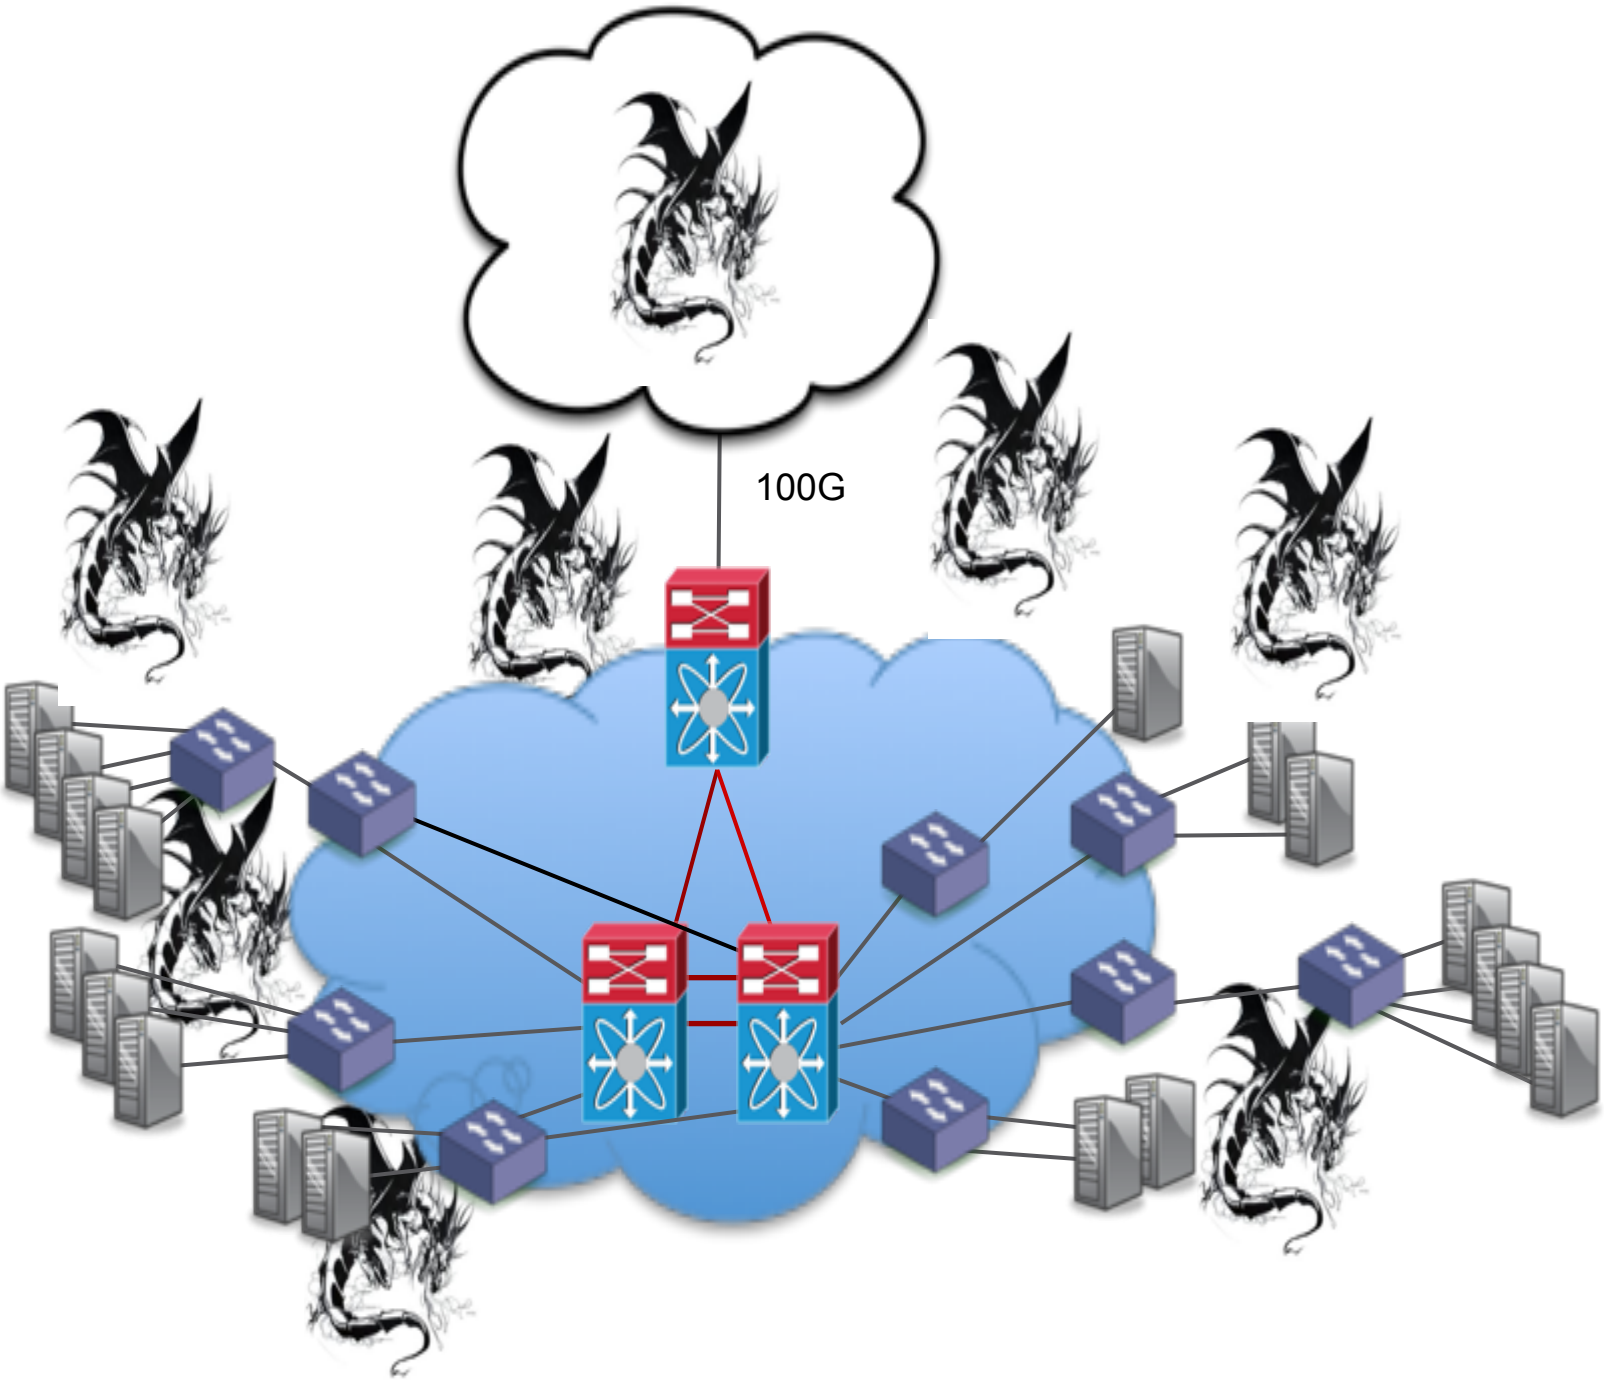
100G

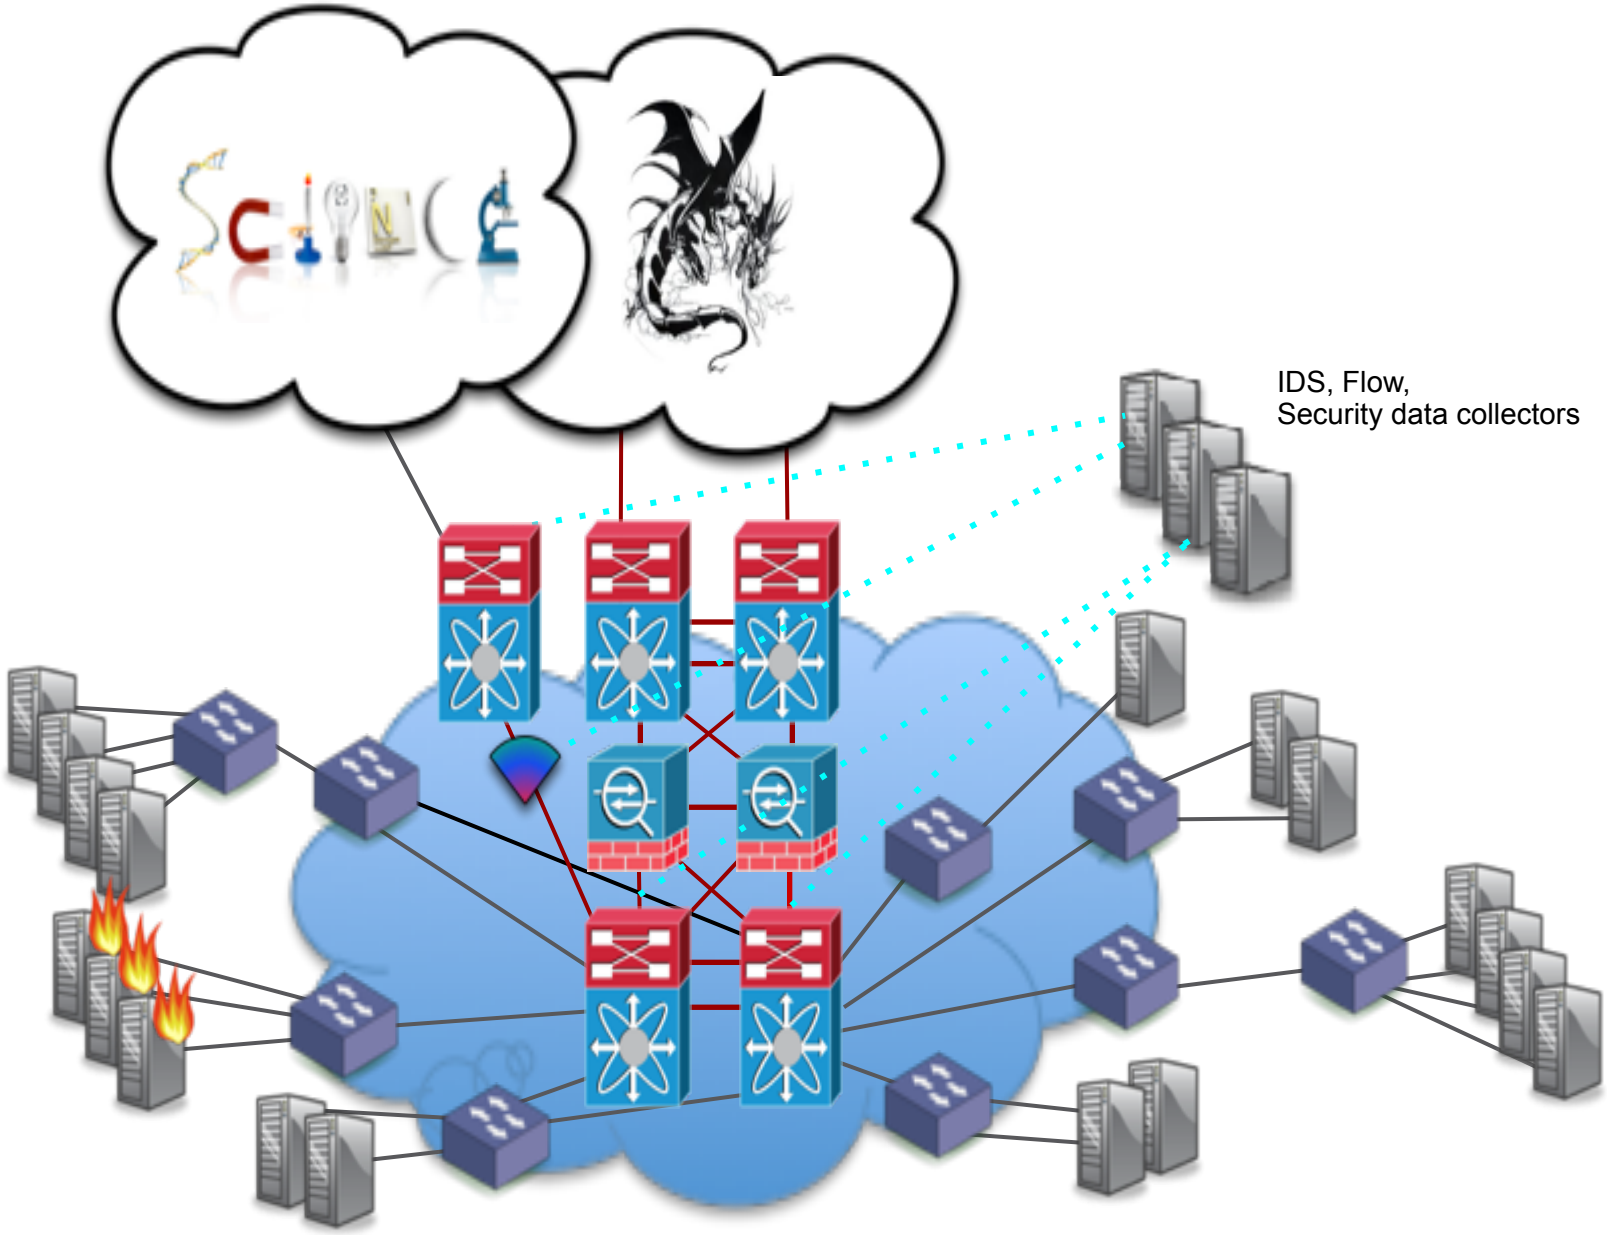
IDS, Flow,
Security data collectors

How does your existing security work?

- Perimeter Security
- Patch Scheduling
- Host integrity
- Data assurance
- Accountability
- Action

Perimeter Access Control

- Best Practice ACLs
  - Block access to control plane
  - Deny inbound access to known exploitable protocols

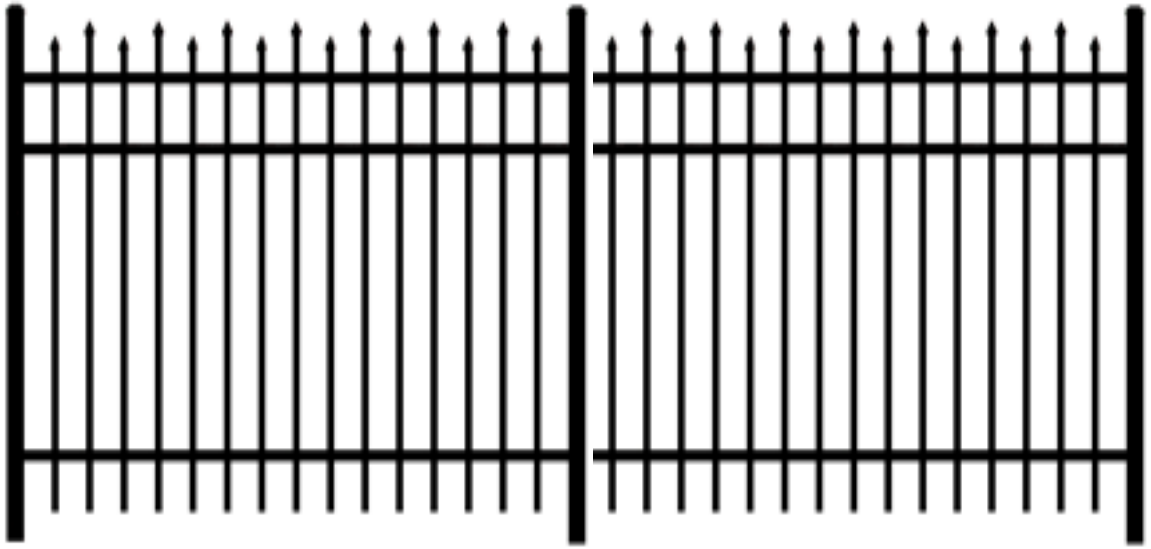

Limit exposure

- Announce only what needs to access research resources
  - Where reasonably possible, announce only research resources via science DMZ

Software Patching

- Patch Scheduling

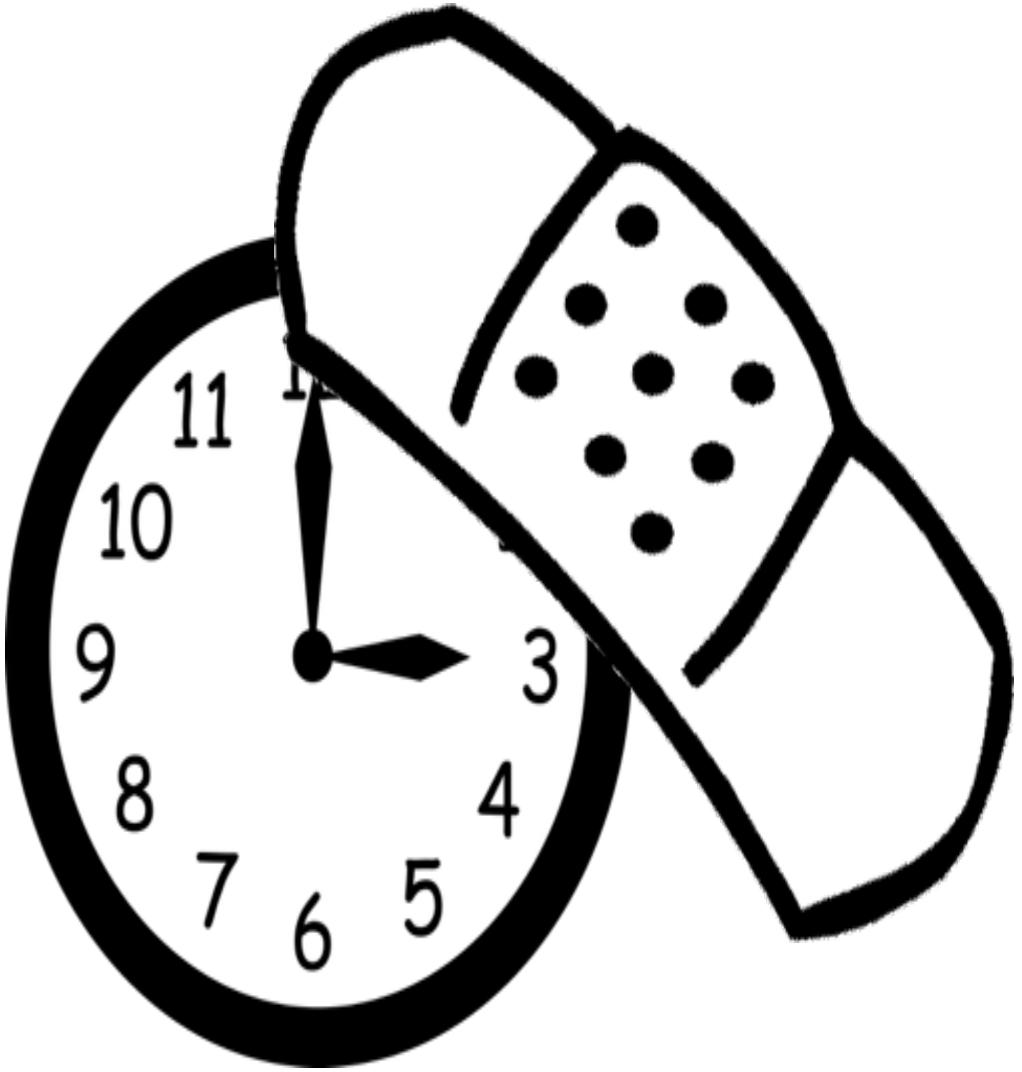

Host Based firewalls

- Host Security - Host based Firewalls

Central Management

- Host Security - Central Management

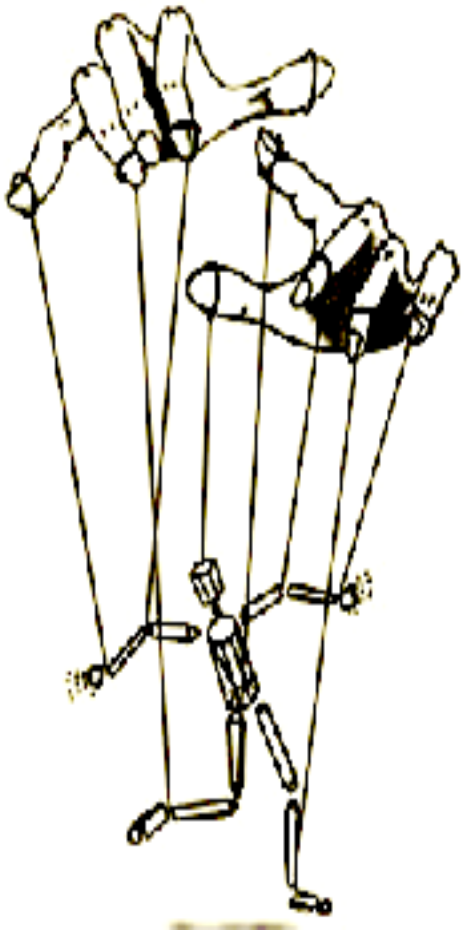

Host IDS

- Host Security - HIDS (Host IDS)

Accountability

- User Accountability

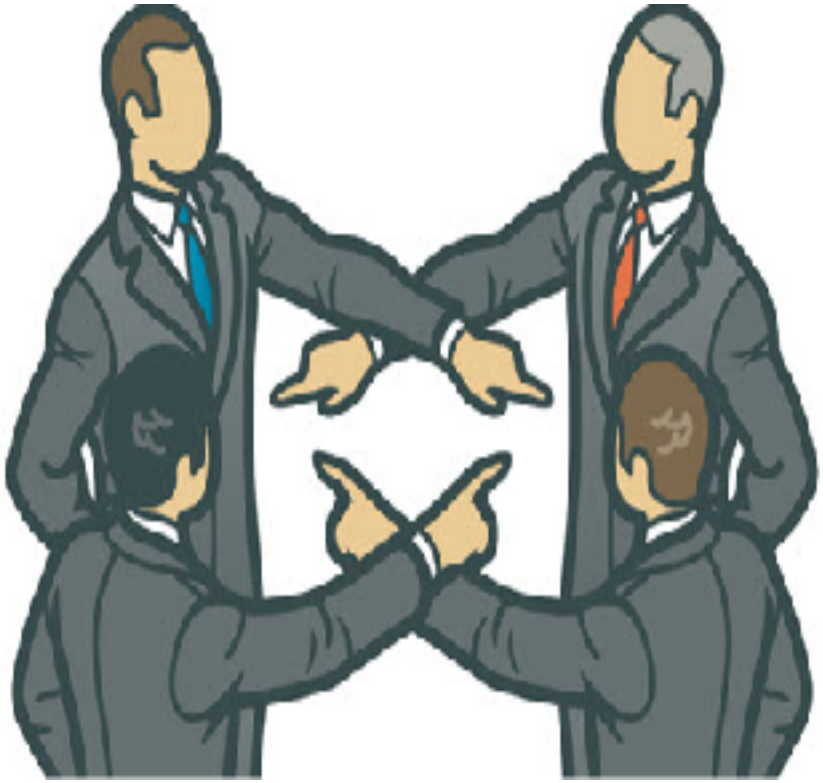

Logging

- Log aggregation

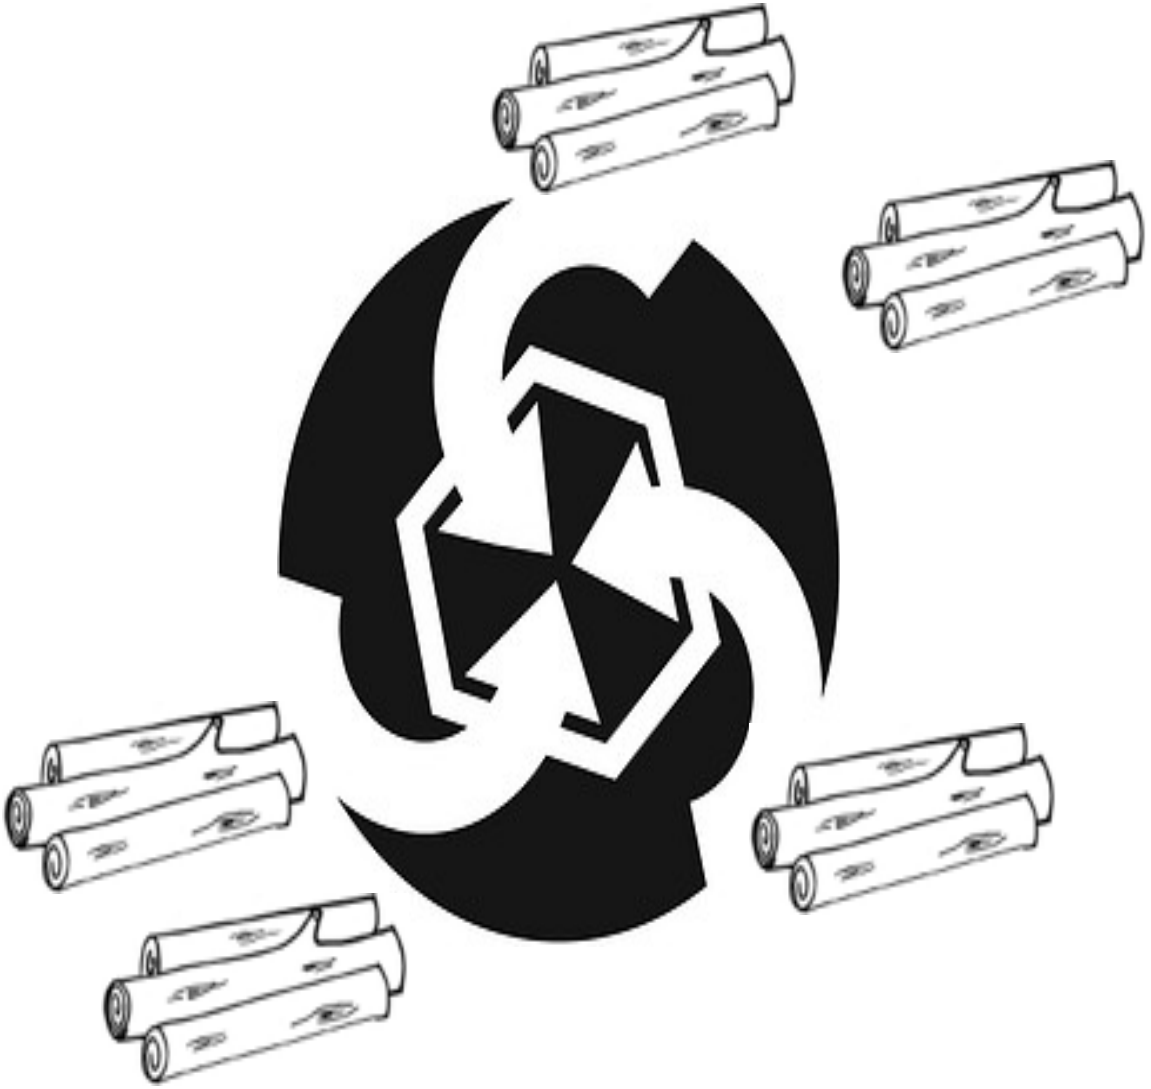

Confidentiality

- Use secure protocols whenever possible
- Utilize SHA2 and other data verification mechanisms

Heavy Lifting

- Intrusion detection system

Action

- Dynamic black hole routing
- BGP FlowSpec (RFC 5575)
- Community feeds (Bogons, etc.)

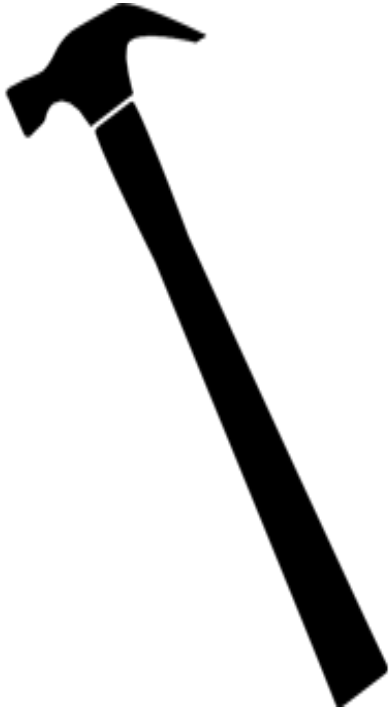

Action – Black Hole Routing

- Dynamic black hole routing
  - Community BGP feeds (Bogons, etc.)

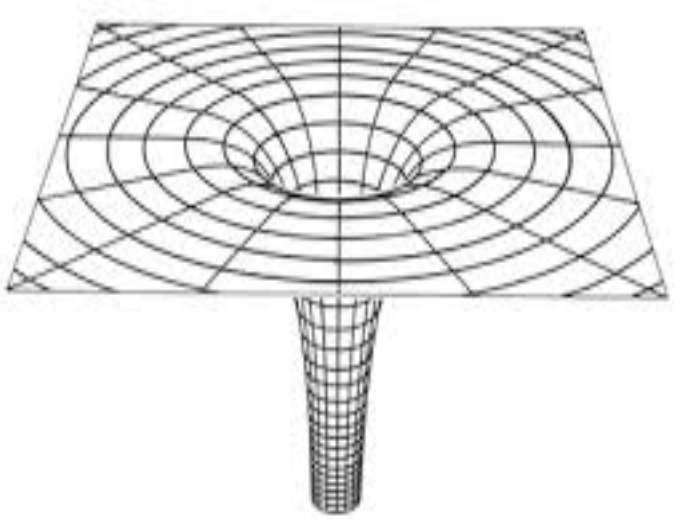

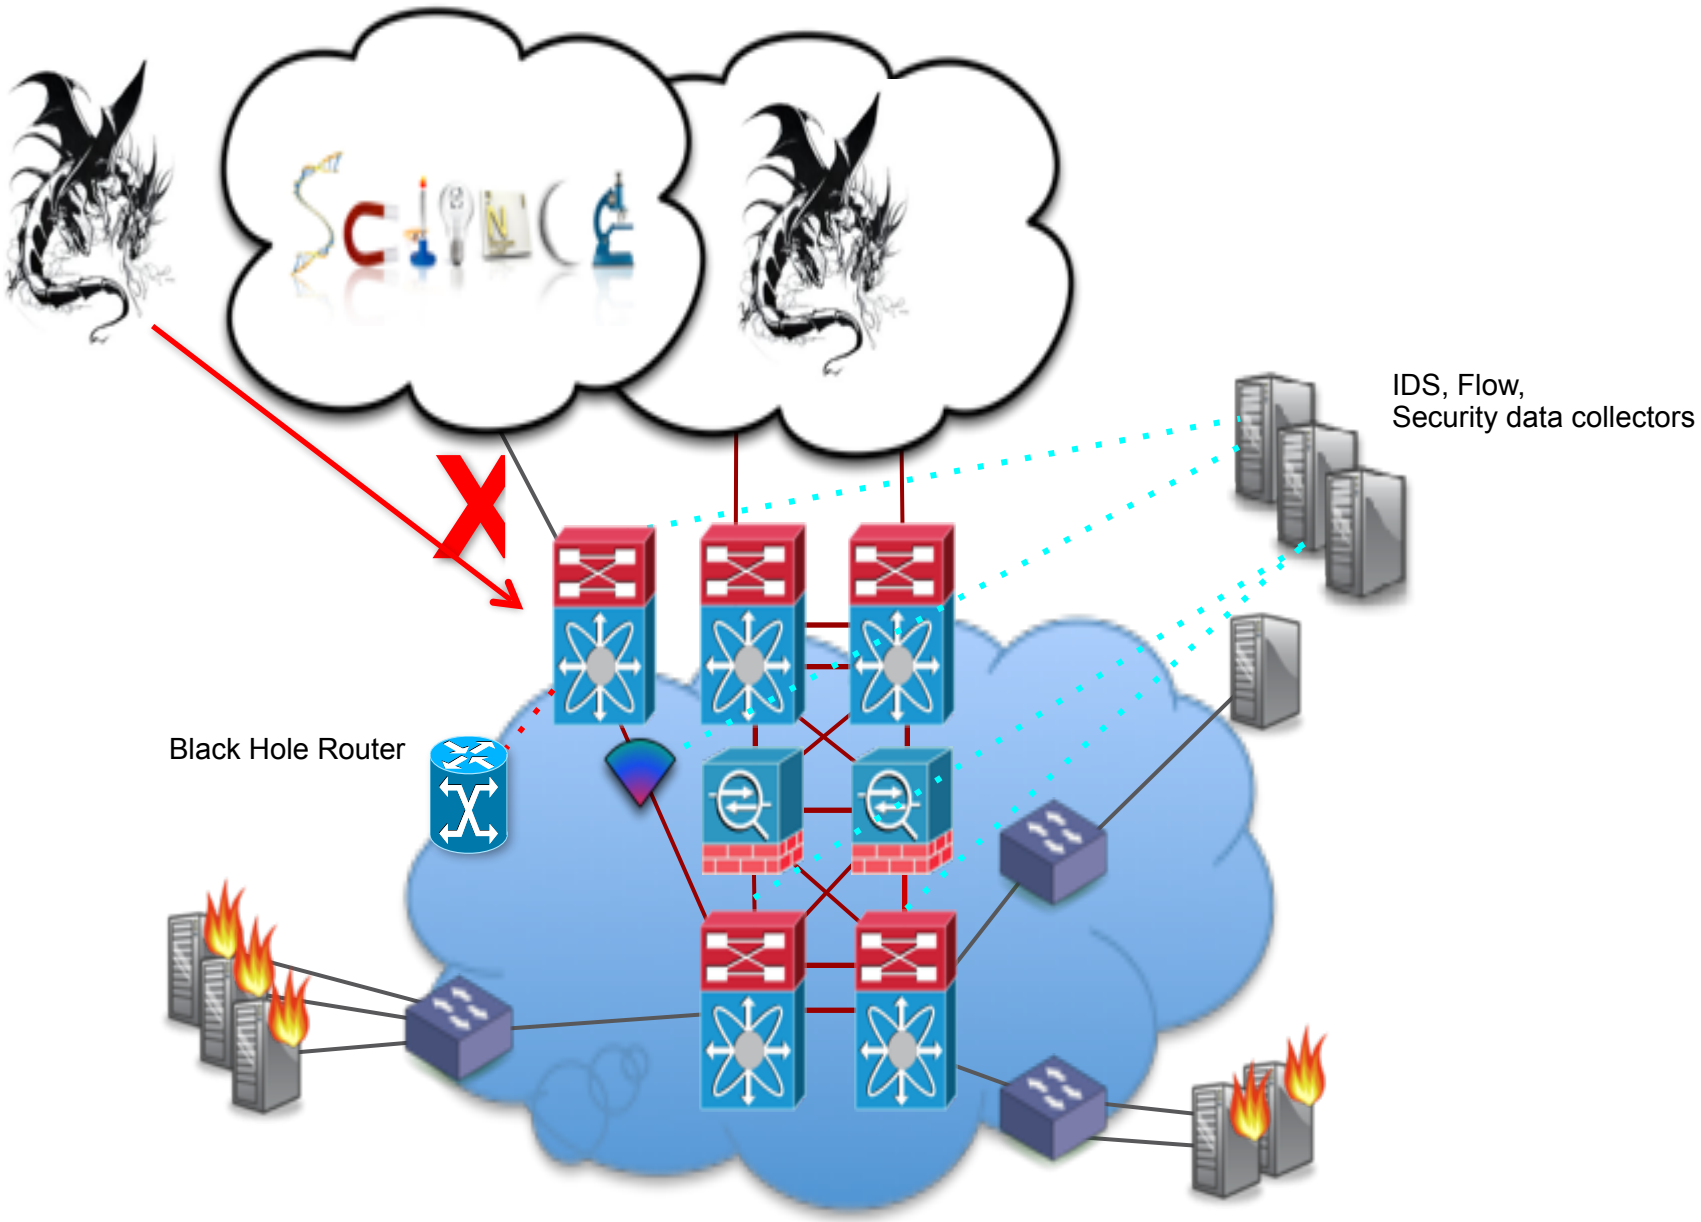
IDS, Flow,
Security data collectors
Black Hole Router

Action – BGP FlowSpec

- Dynamic black hole routing
  - Dissemination of rules via BGP NLRI

RFC 5575

Baselines

- Traffic graphs
- Flow Data
- Syslog (host and network)

IPv6

- Don’t forget IPv6

Notable mentions

- SDN

Collaboration

- Multiple groups working together

Useful tools and Links

- engage@es.net
- http://fasterdata.es.net/science-dmz/science-dmz-security/
- http://www.bro-ids.org

Example Checklist

- Announce only what needs to access research resources
- ACL control plane services of all network, storage and management hardware
- Host based firewalls
- Central host management service
- Central syslog
- Flow data
- SNMP counters and graphs
- Regularly scheduled external vulnerability scanning

Science DMZ makes your network more secure

- We have talked about how to make the Science DMZ more secure. Now, how do we make your network more secure using good Science DMZ practices?
- Scenario 1: Scientific instruments
- Scenario 2: High-performance computing

Scenario 1: Scientific Instruments

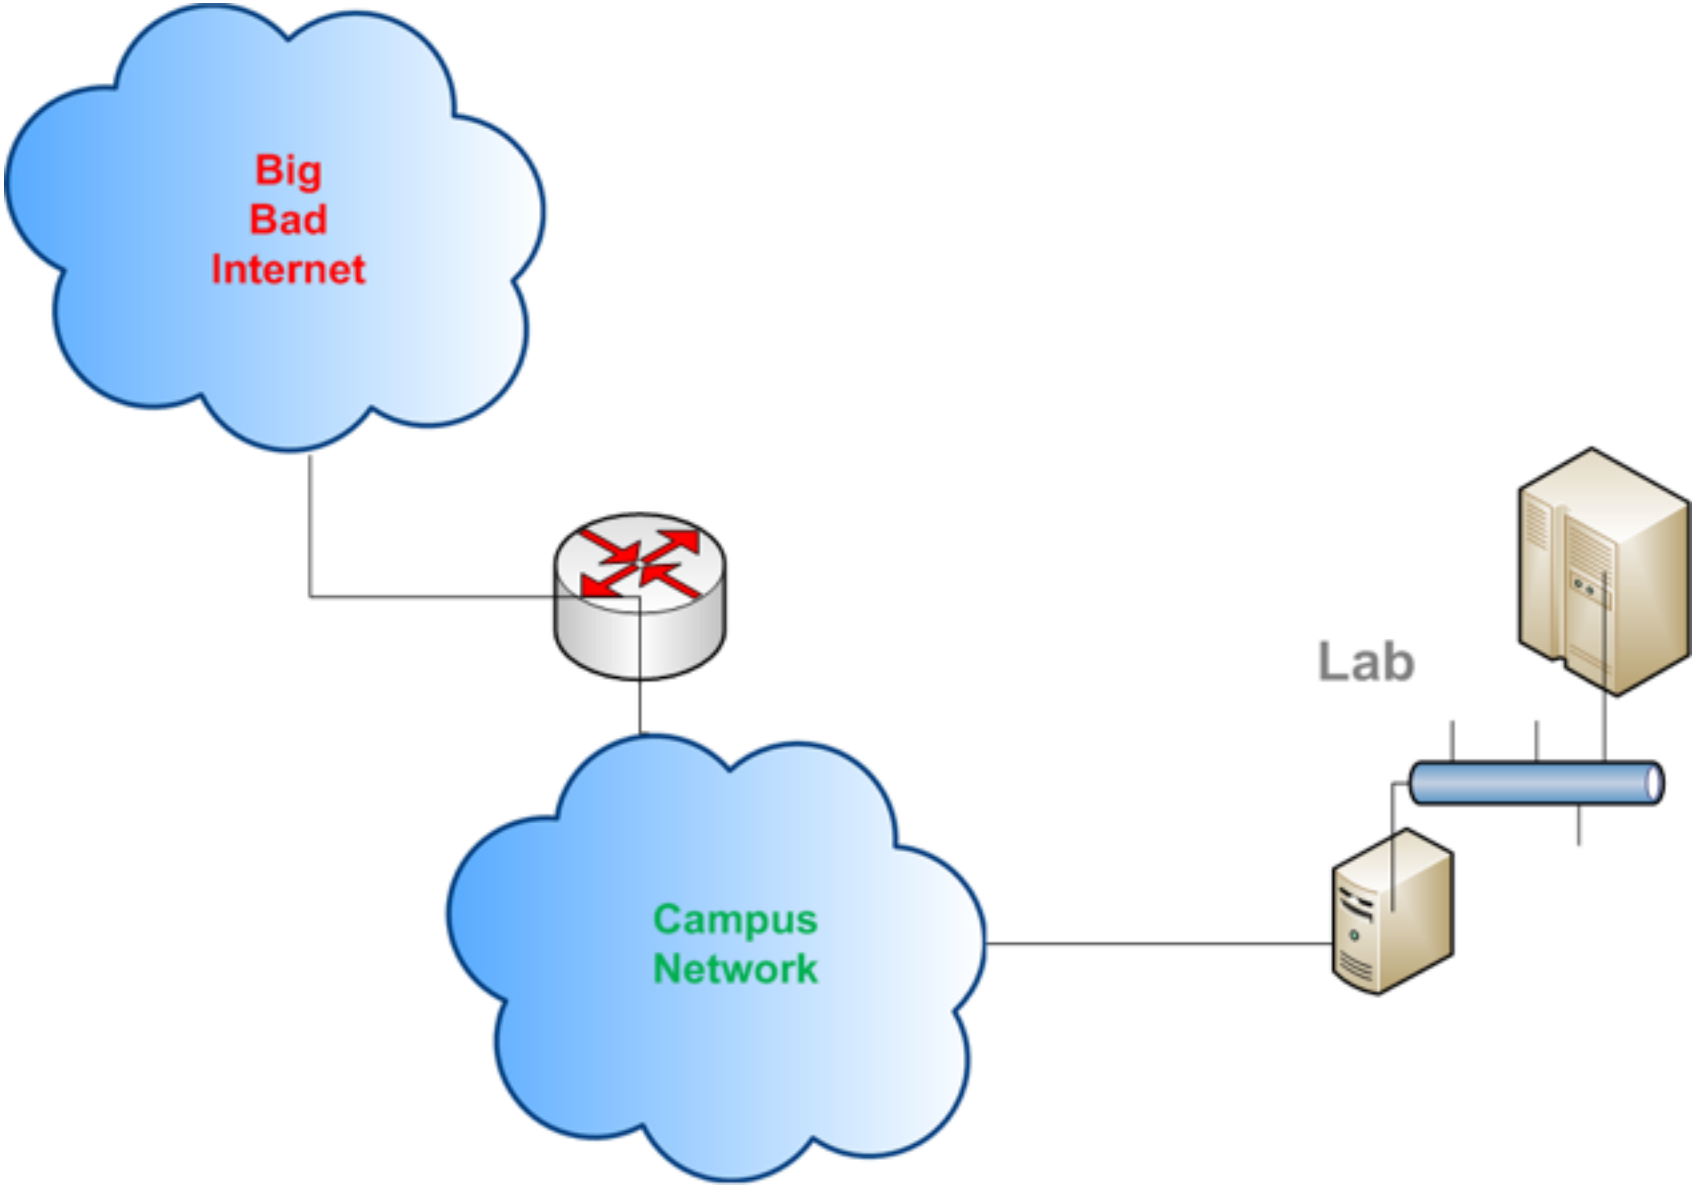

Scenario 1: Scientific Instruments

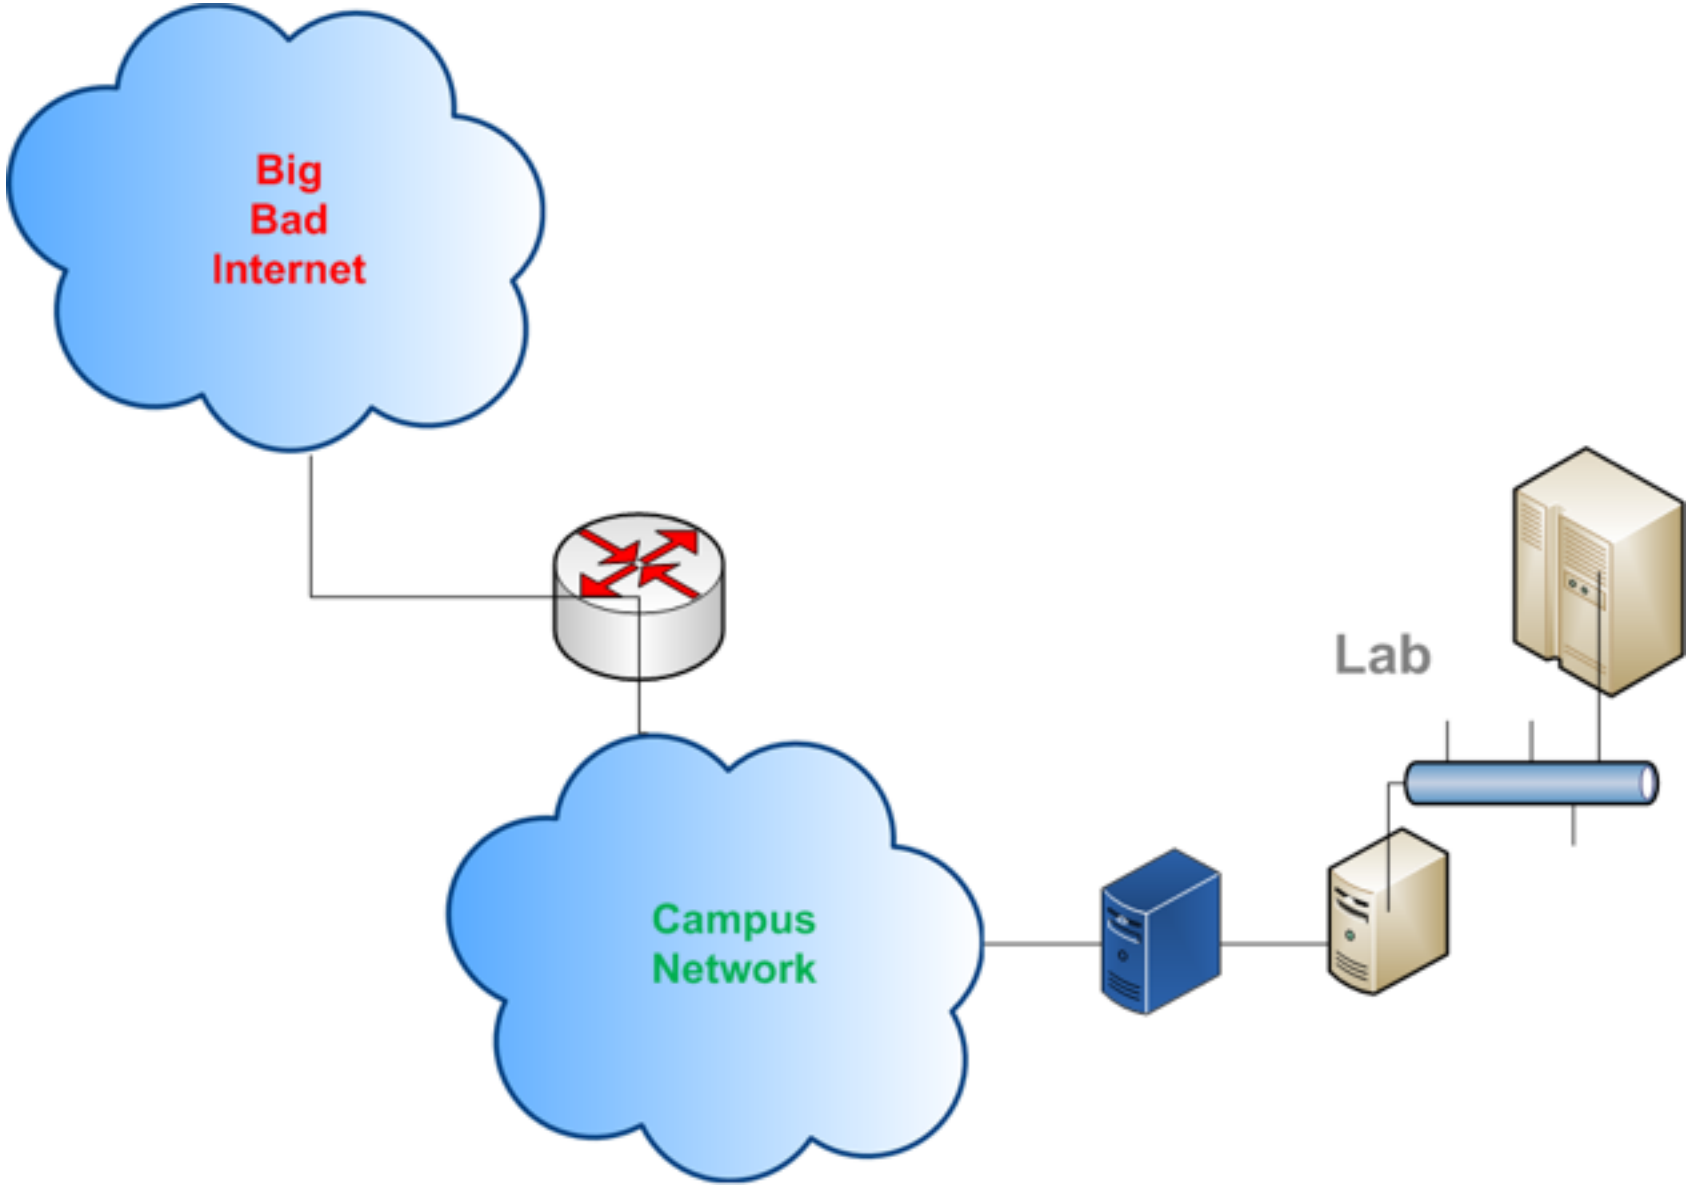

Scenario 1: Scientific Instruments

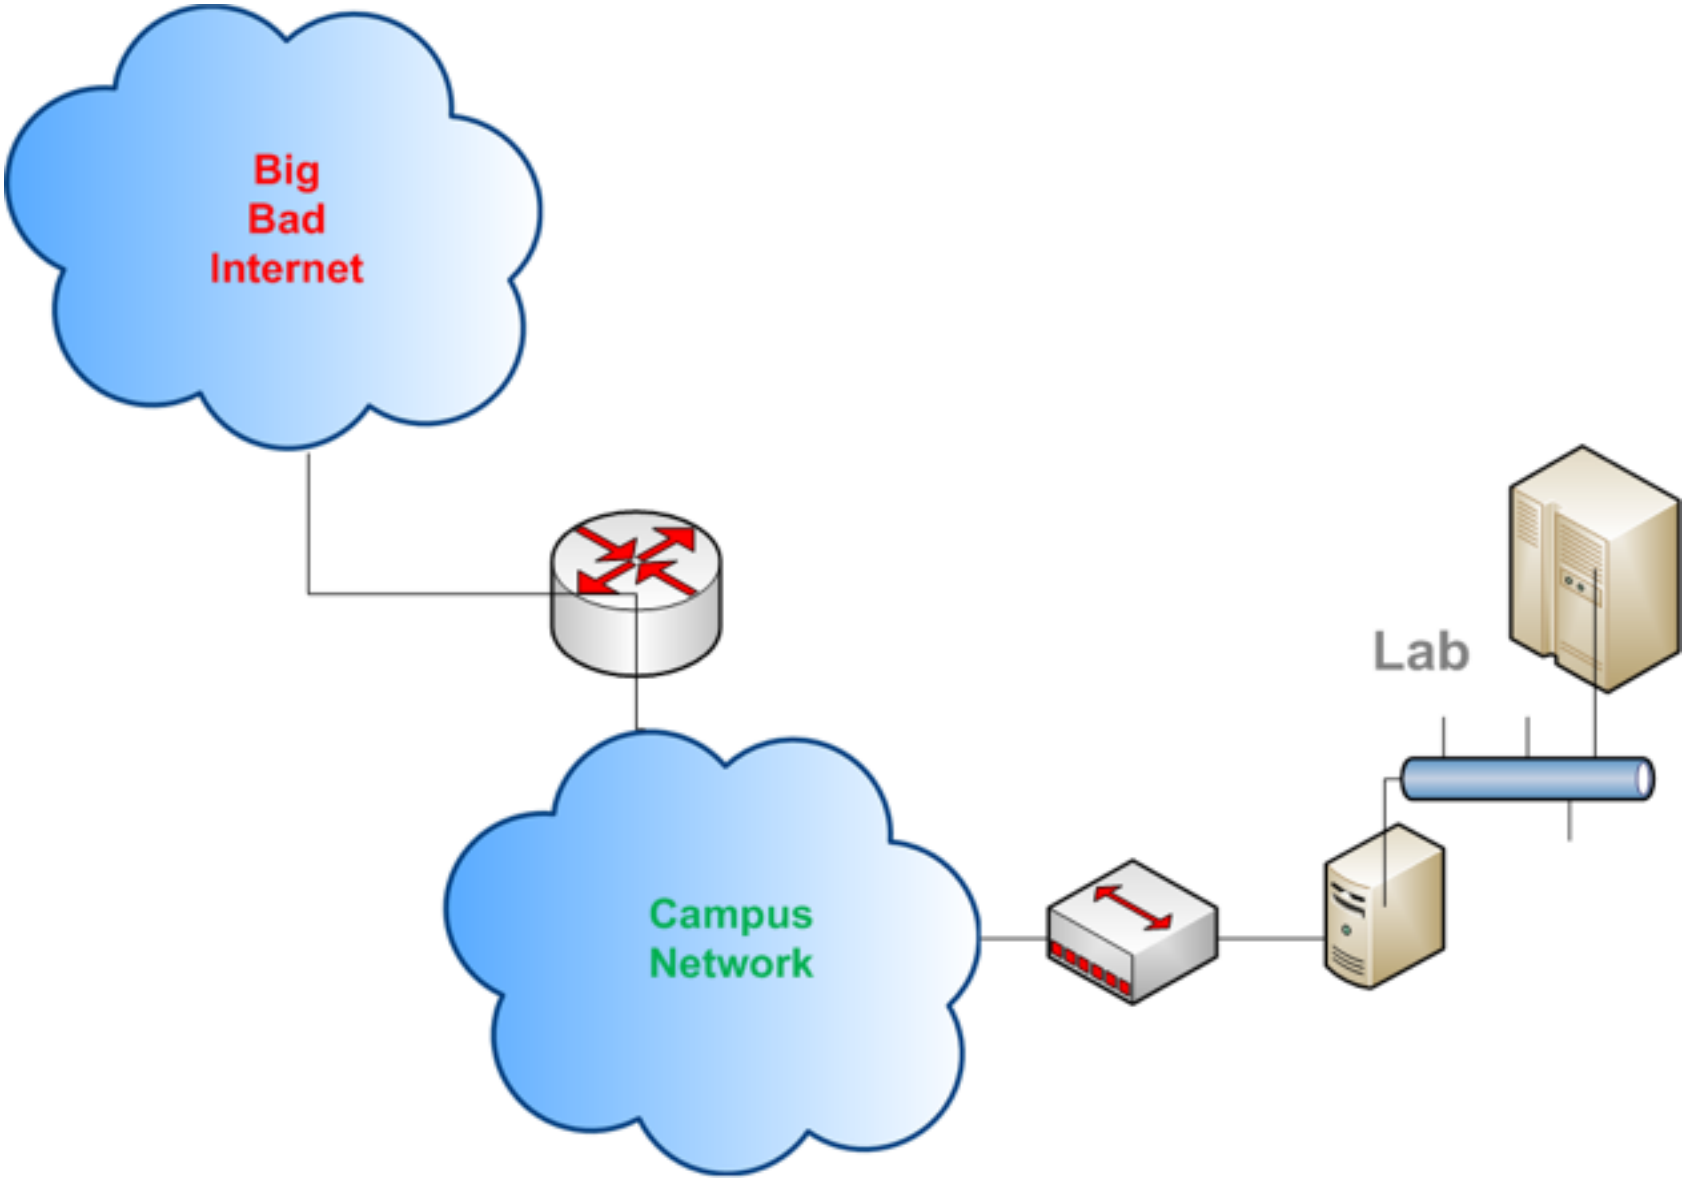

Scenario 1: Scientific Instruments

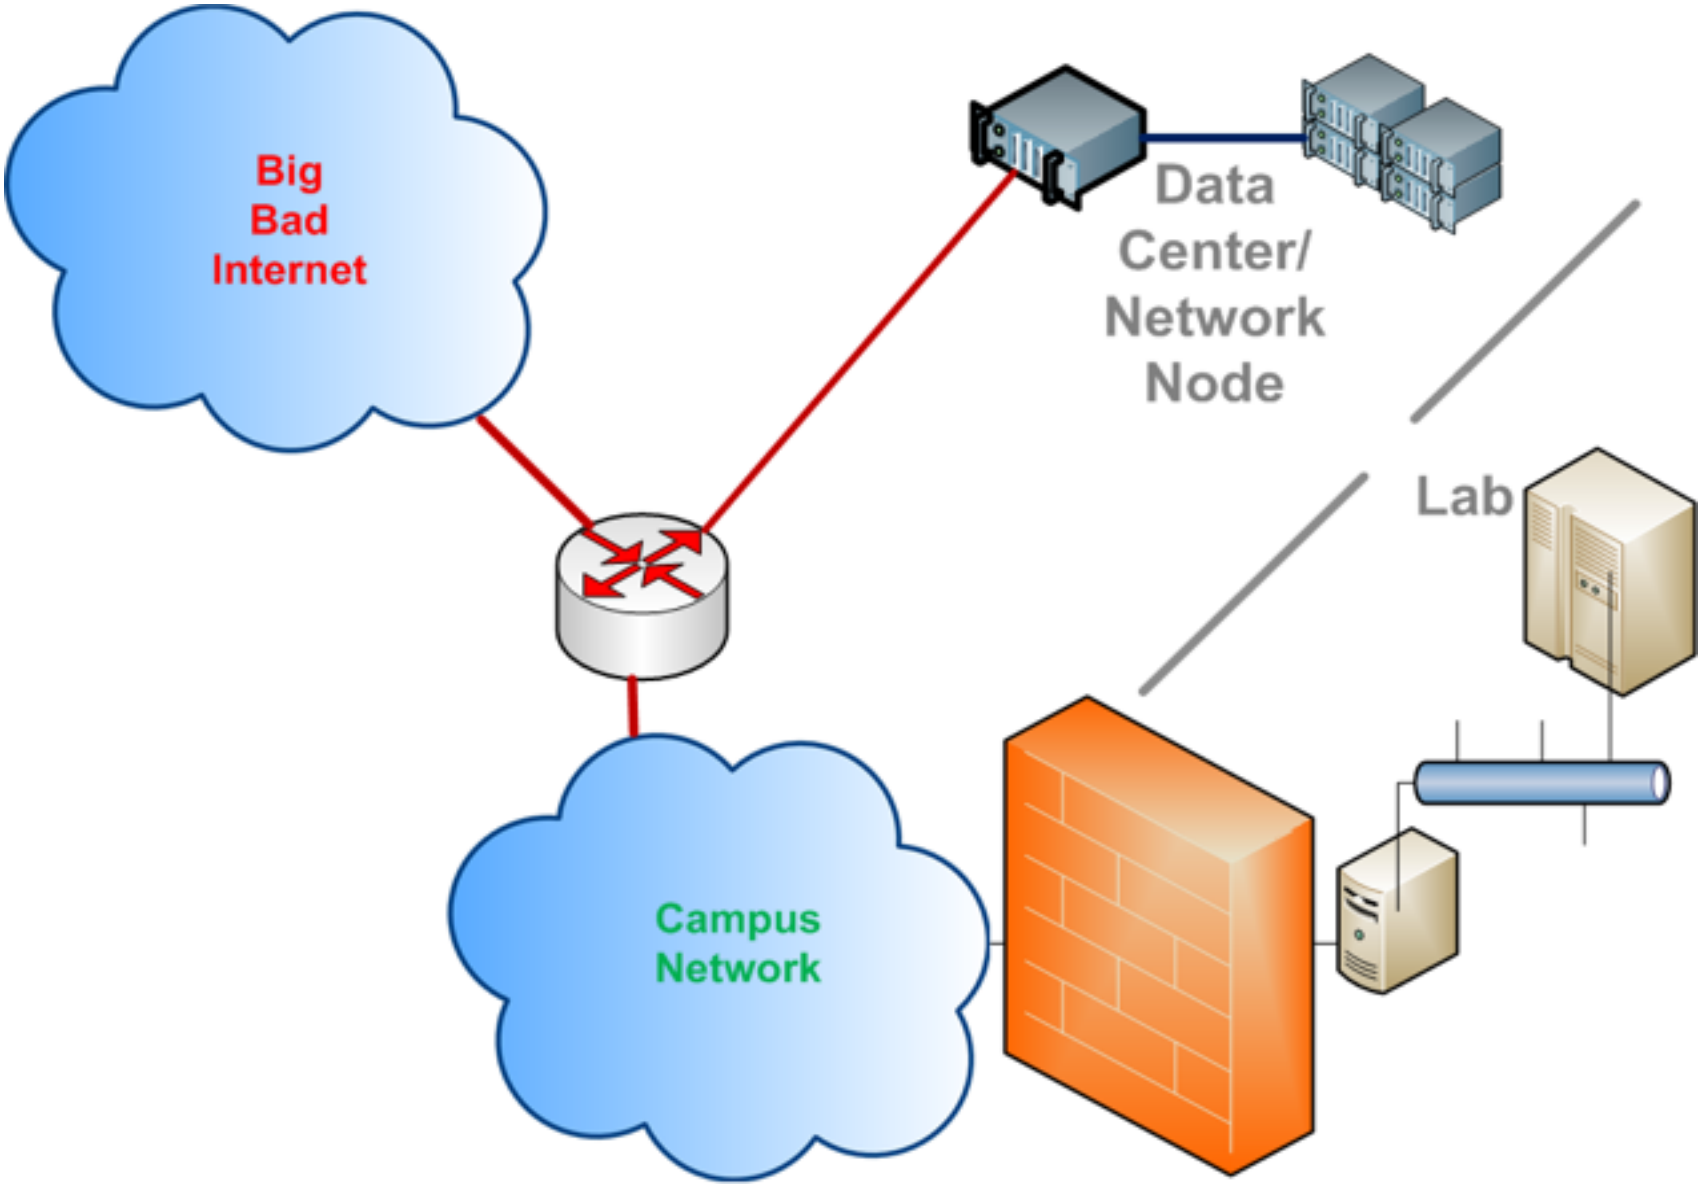

Scenario 1: Scientific Instruments

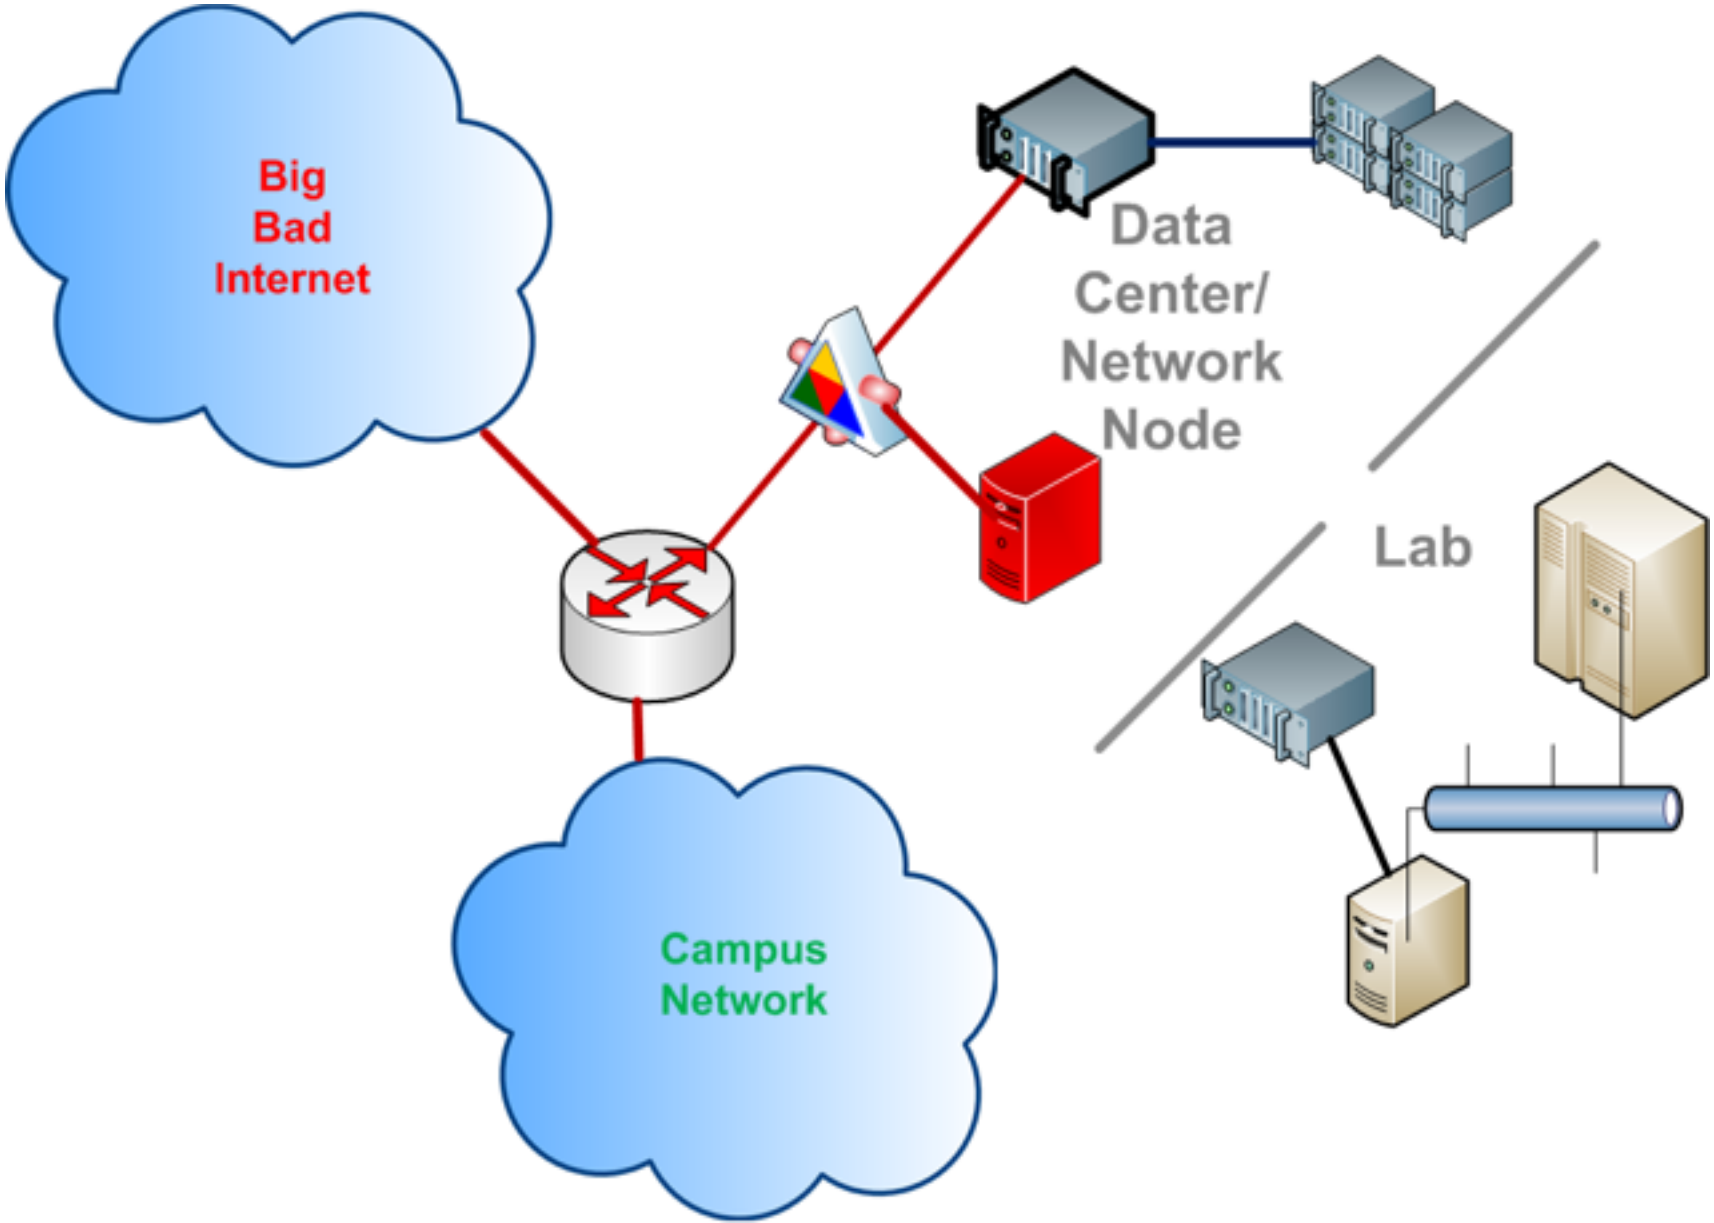

Scenario 1: Scientific Instruments

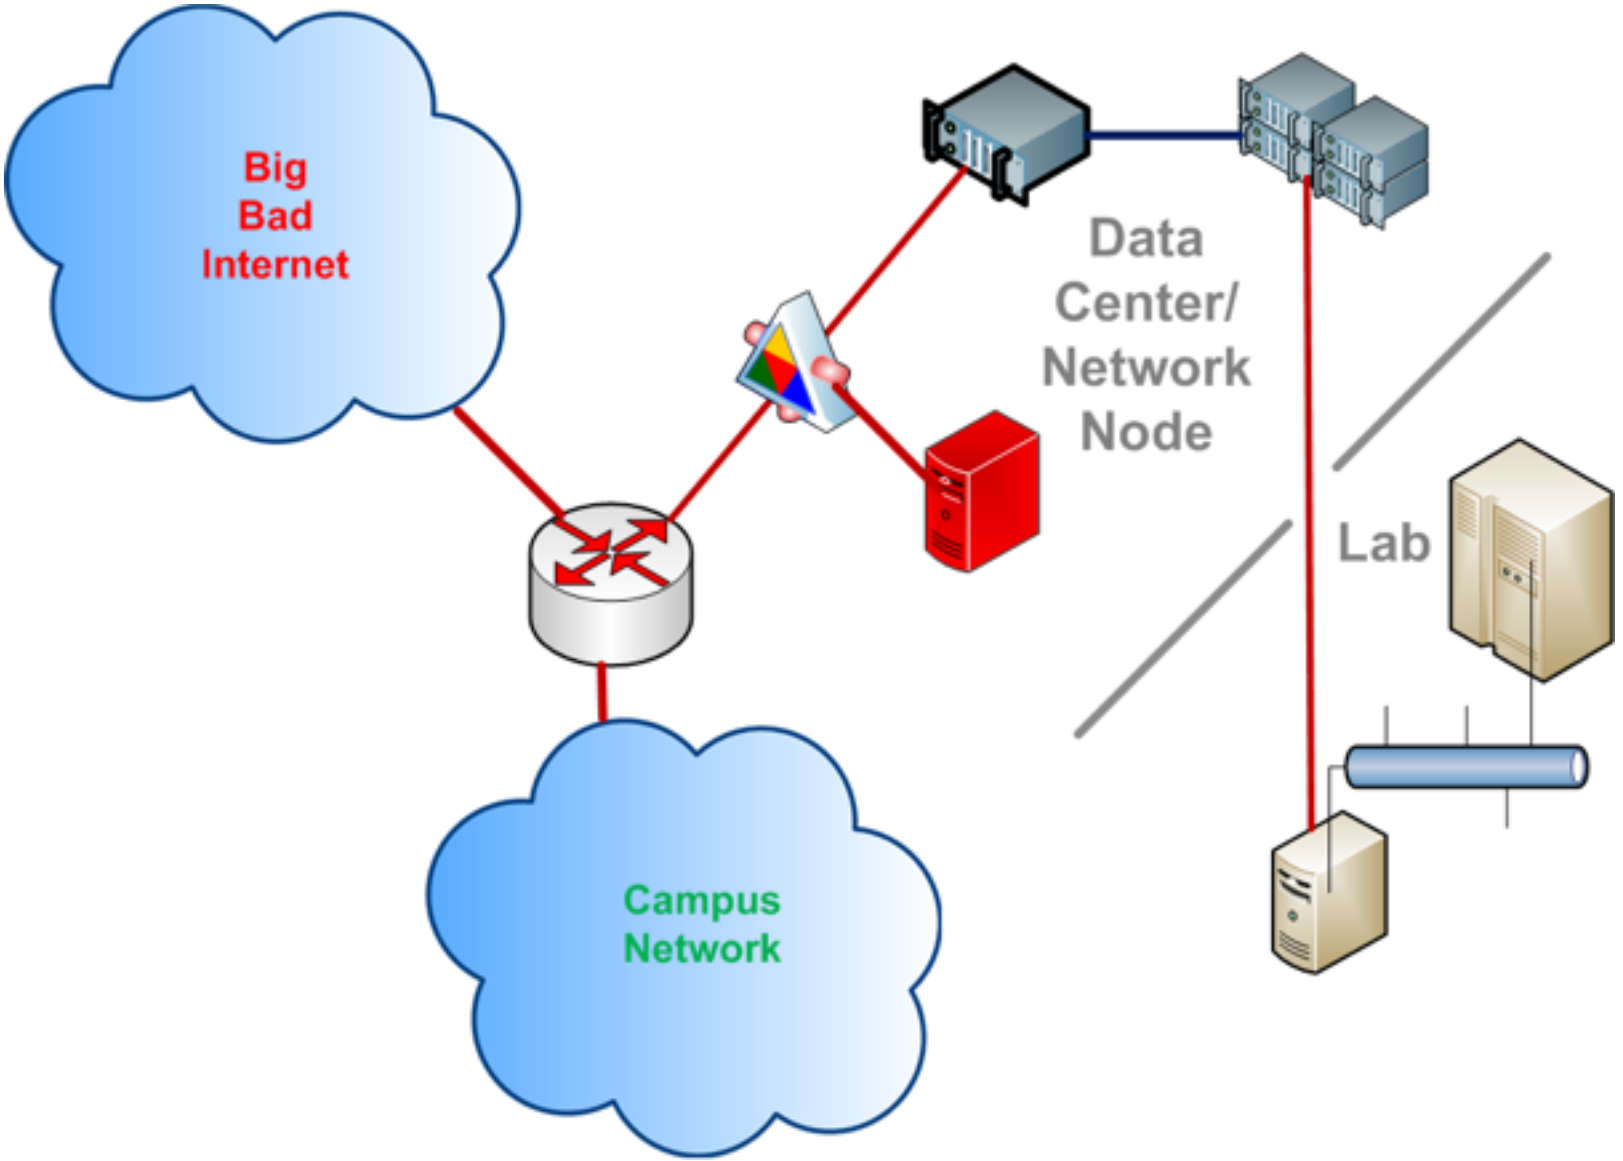

Scenario 1: Scientific Instruments

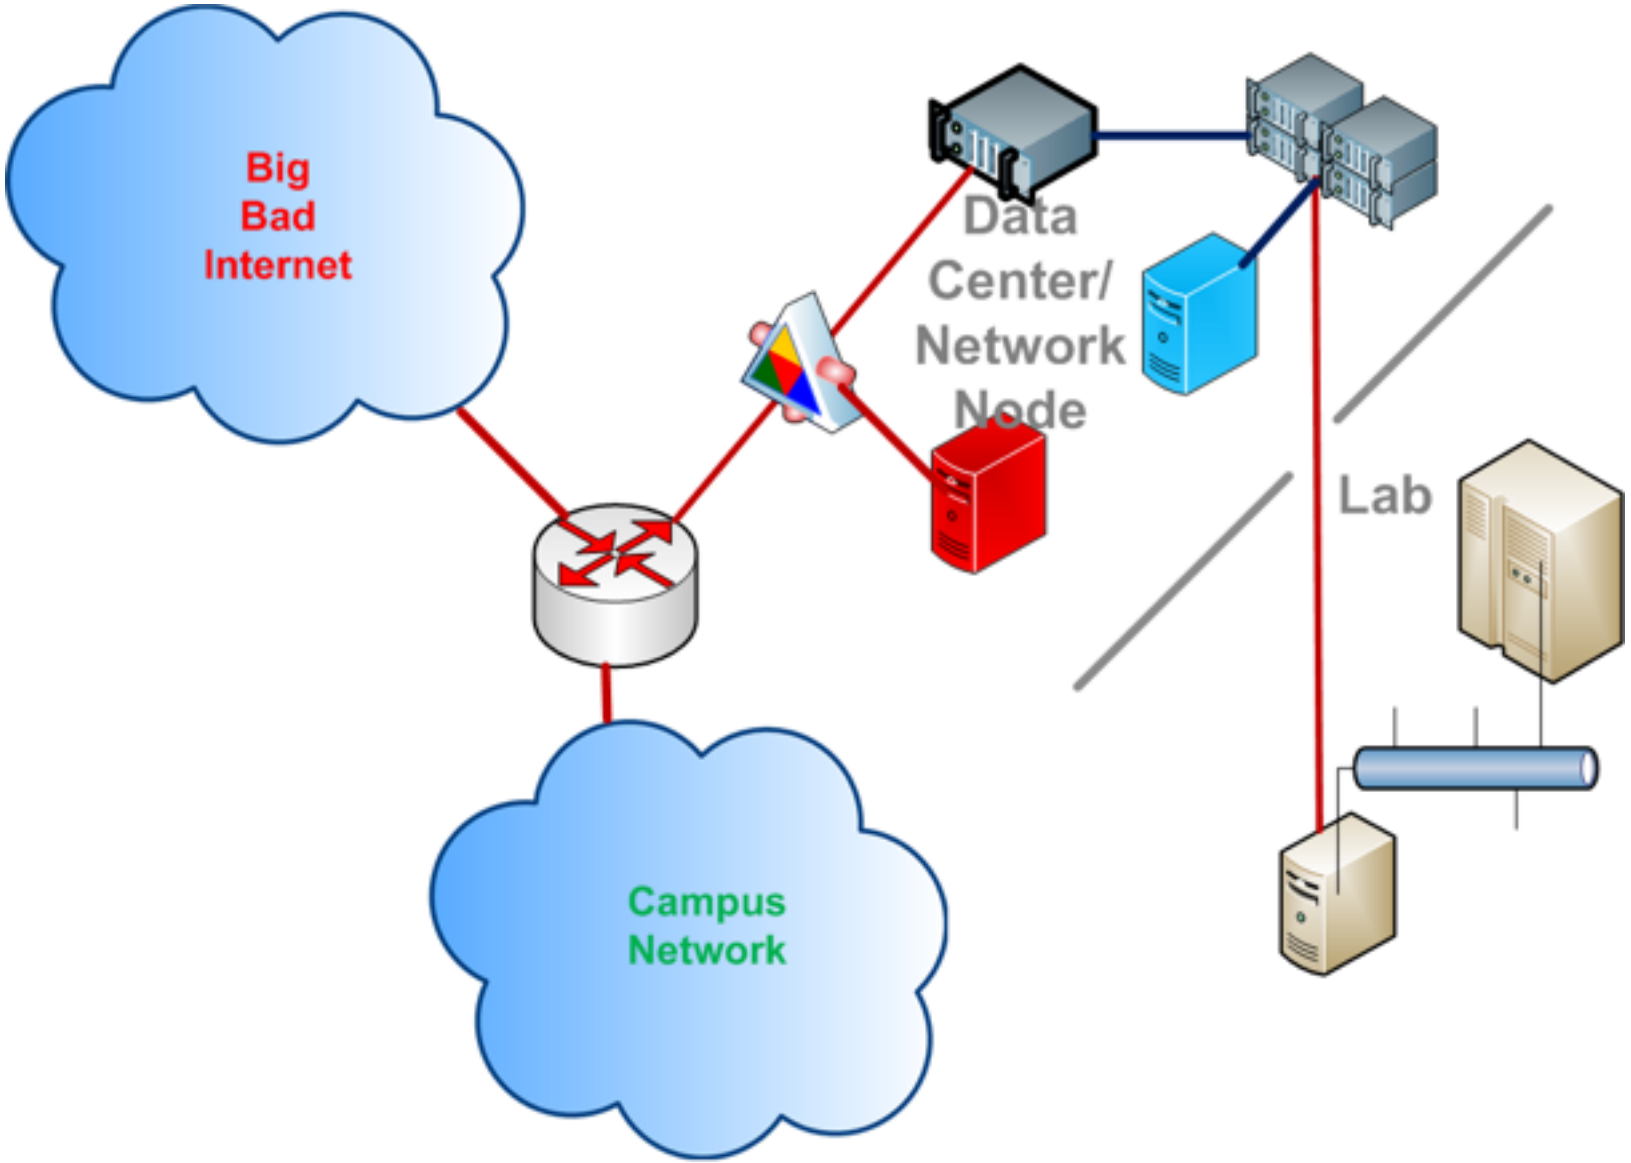

Scenario 1: Scientific Instruments

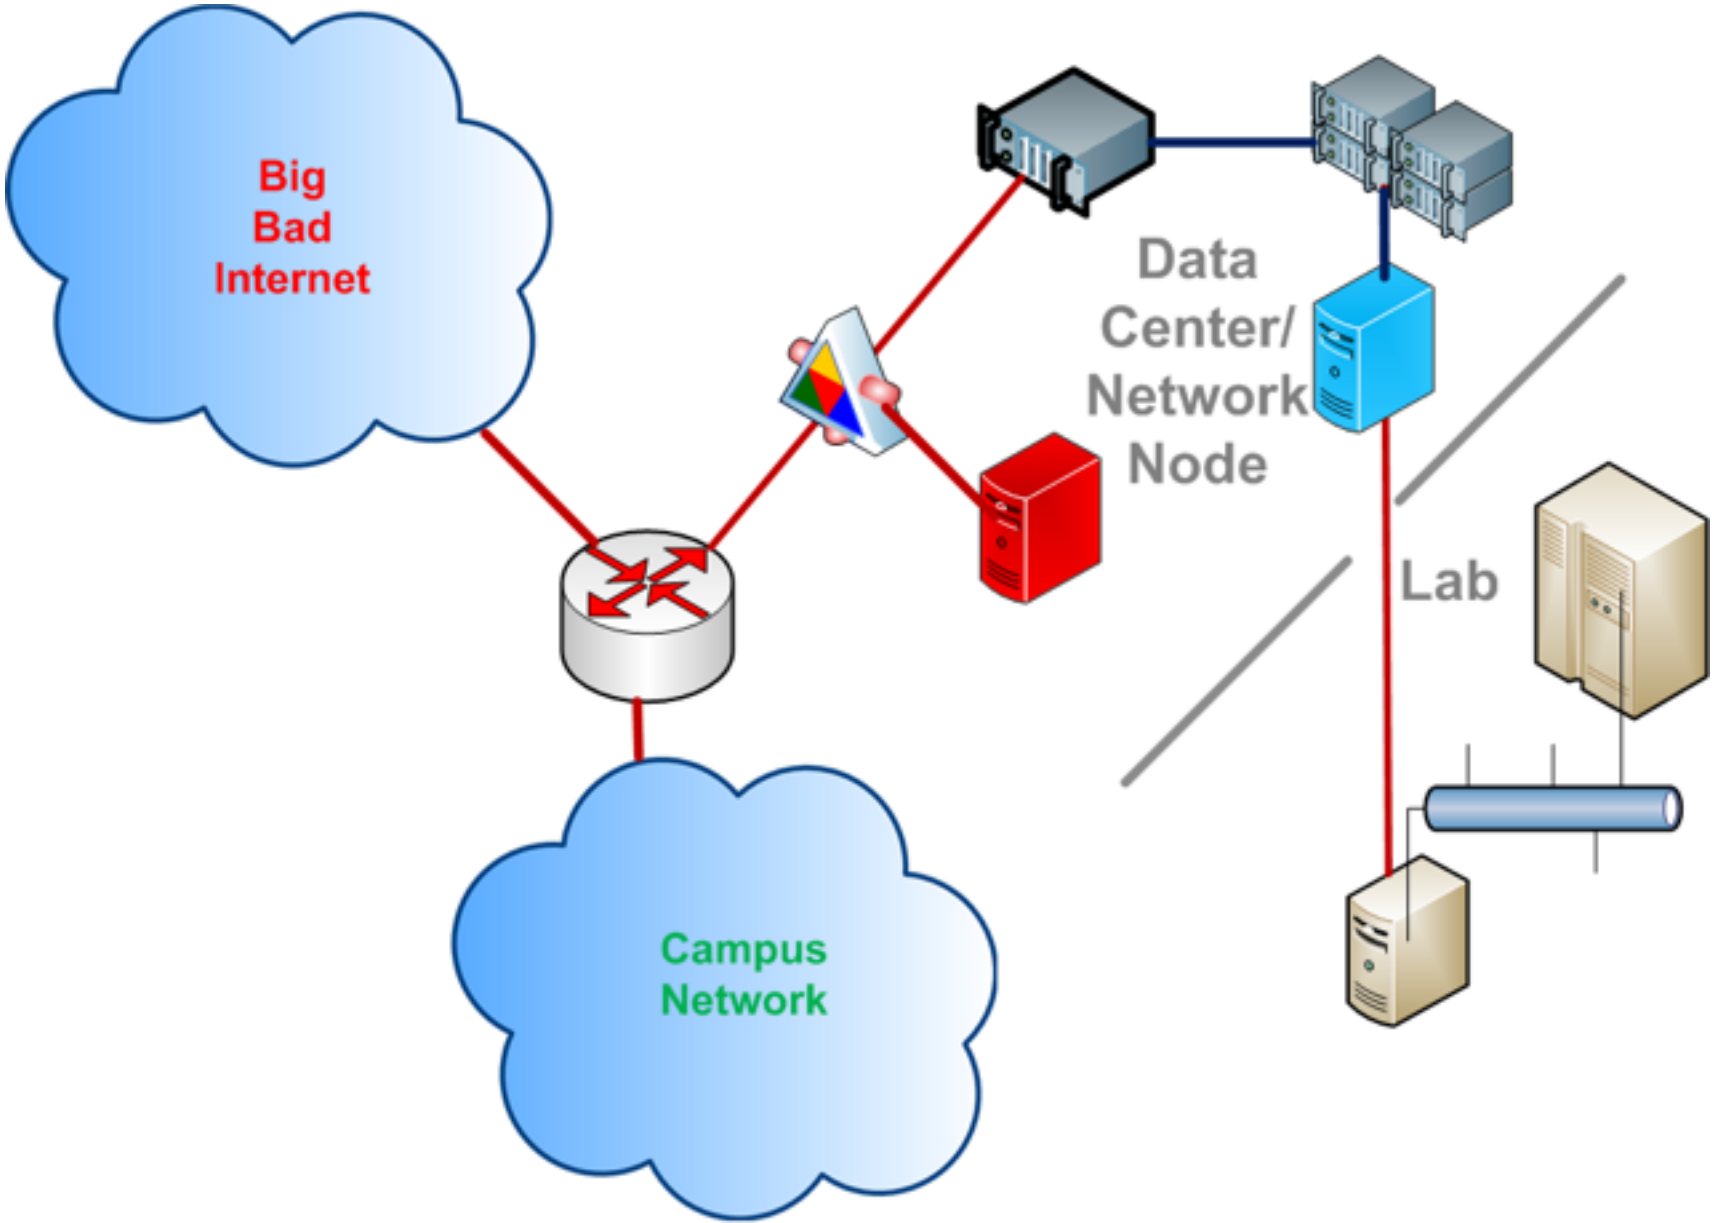

Scenario 1: Scientific Instruments

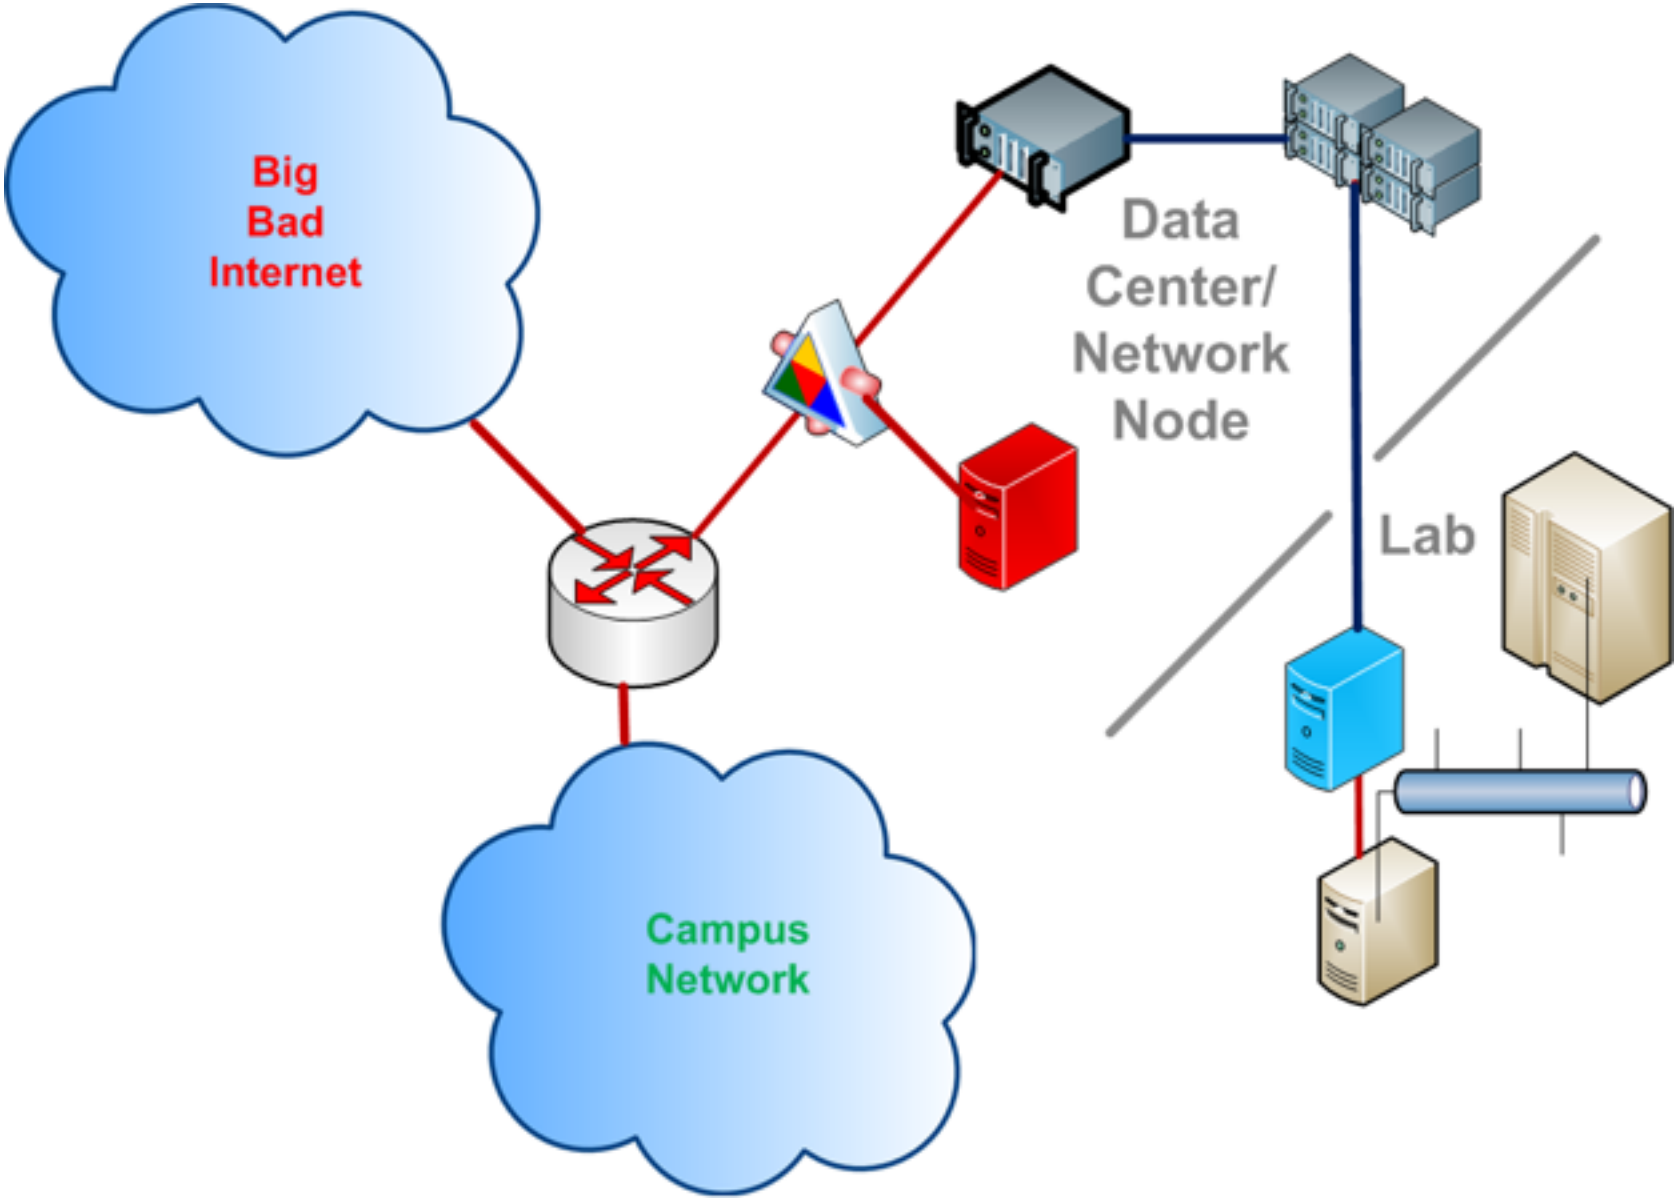

Scenario 2: Compute Clusters

- Compute clusters may have specialized software for scheduling jobs or managing parallel nodes and resources.
- Most nodes may be on private network.
- Bastion hosts, with various AUTHNZ schemes – may also need specialized software:
  - 2FA
  - Instrumented SSH
- DTNs may also need specialized software:
  - Globus
  - High-throughput data transfers
  - Special filesystems

Scenario 2: Compute Clusters

- In such a situation, your compute cluster should not also be your DTN.
- Much easier to secure if you separate these functions.
- Try to keep things as standard as possible on as many machines as possible.
- Separation of functions allows for better risk-assessment and more carefully-tailored controls.
- Controls should be matched to the thing that you're protecting.
- Avoid one-offs if possible, but if you have to have them, make sure they're well-designed, well-managed, and well-documented!
- The Science DMZ helps with all of these things.

Conclusions

- Separation of functions (resulting in good network segmentation) is key.
- The Science DMZ makes this possible.
- A well-designed Science DMZ is a security architecture.
- The Science DMZ *improves* risk mitigation.
- The Science DMZ is not a security workaround. A secure Science DMZ is security.

Questions? (I.e. Do I really need a slide with a question mark to get you to ask me questions?)