1986–2016
ESnet
30 YEARS OF NETWORKING
AT THE SPEED OF SCIENCE

The Science DMZ as a Security Architecture

Michael Sinatra
Network, Systems, Security Engineer
Energy Sciences Network
Lawrence Berkeley National Laboratory

NSF Cybersecurity Summit for Large Facilities and Cyberinfrastructure
Arlington VA
17 August 2016

U.S. DEPARTMENT OF ENERGY
Office of Science

BERKELEY LAB

Who am I? Why am I here?

- Served on several security committees and "big incident" response teams at UCB.
- Limited time security strategist for ESnet.
- Worked with Nick Buraglio within ESnet to develop security controls tailored to the Science DMZ.
- Interested in Science DMZ for many years…

Motivations

- I have more recently been a bit concerned about how security is “done” in R&E.
  - Too much top-down policy and “control” orientation. (This was necessary at one point, but I am not sure it is now.)
  - Checkbox compliance.
  - Lack of good risk assessment.
  - Failure to account for network functional needs (leading to Joe St. Sauver’s idea of a “Network Usability Officer).
  - Equating “controls” with “security.”
- The Science DMZ has emerged out of a similar set of concerns, but we’re currently hampered by some myths.

Motivations

- The big myth: The main goal of the Science DMZ is to avoid firewalls and other security controls.
  - Leads to all sorts of odd (and wrong) claims like:
    - "Our whole backbone is a Science DMZ because there is no firewall in front of the backbone."
    - "The Science DMZ doesn't allow for **any** security controls."
    - "The Science DMZ requires a default-permit policy."
  - The reality is that the Science DMZ emphasizes reducing degrees-of-freedom, reducing the number of network devices (including middleboxes) in the path, eliminating devices that can't perform, and ensuring that the devices that remain in the path are capable of large-scale data-transfer caliber performance.

Motivations

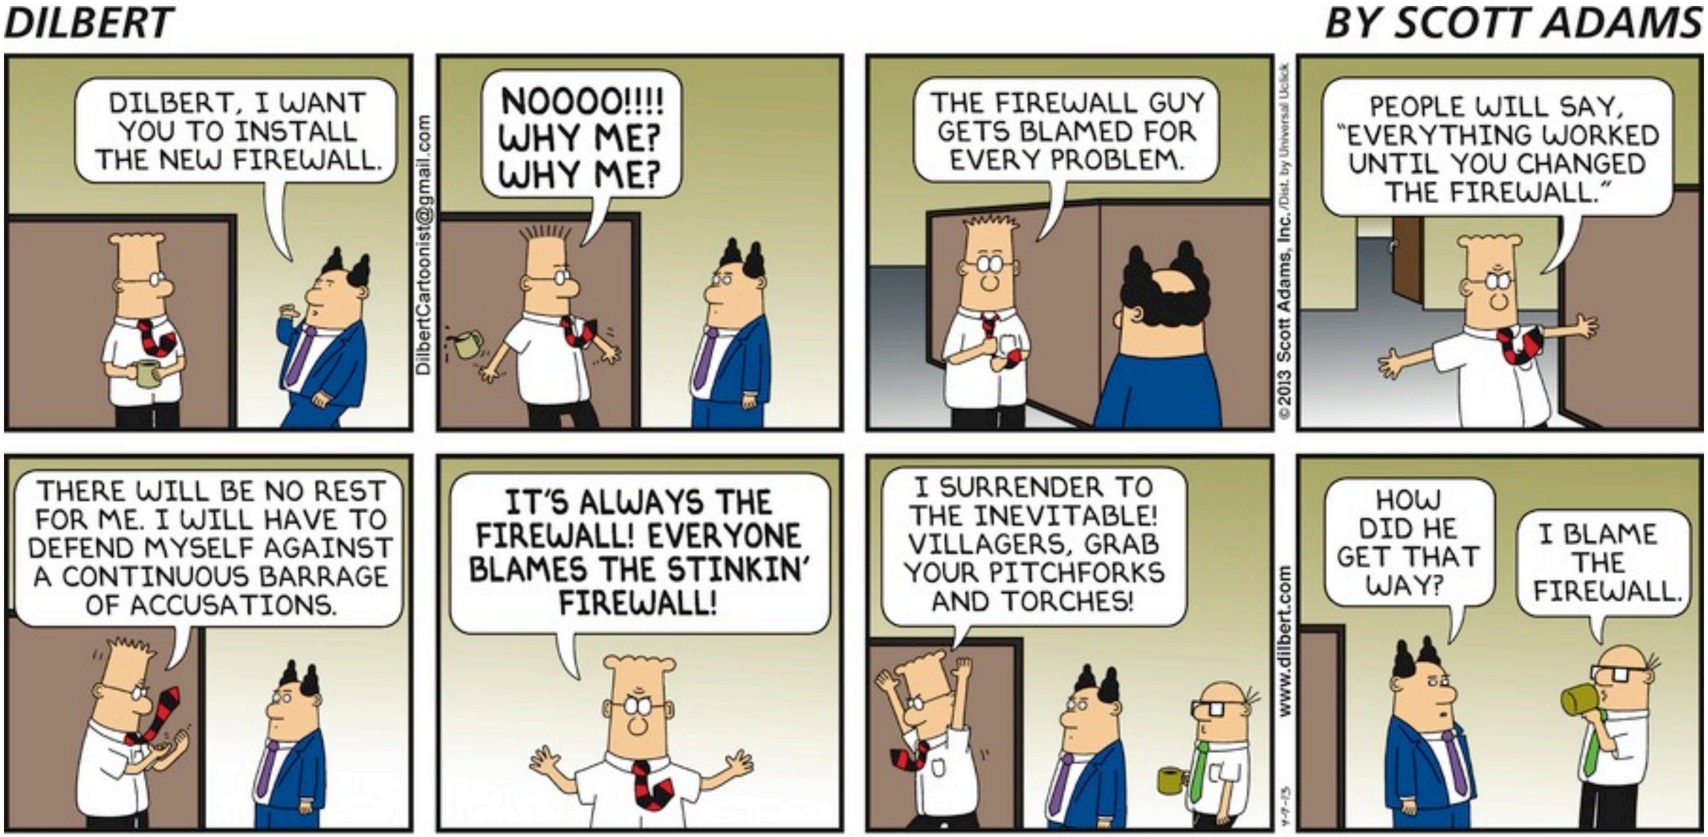
DILBERT
BY SCOTT ADAMS
DILBERT, I WANT YOU TO INSTALL THE NEW FIREWALL.
NOOOO!!!! WHY ME? WHY ME?
THE FIREWALL GUY GETS BLAMED FOR EVERY PROBLEM.
PEOPLE WILL SAY, "EVERYTHING WORKED UNTIL YOU CHANGED THE FIREWALL."
THERE WILL BE NO REST FOR ME. I WILL HAVE TO DEFEND MYSELF AGAINST A CONTINUOUS BARRAGE OF ACCUSATIONS.
IT'S ALWAYS THE FIREWALL! EVERYONE BLAMES THE STINKIN' FIREWALL!
I SURRENDER TO THE INEVITABLE! VILLAGERS, GRAB YOUR PITCHFORKS AND TORCHES!
HOW DID HE GET THAT WAY?
I BLAME THE FIREWALL.
DilbertCartoonist@gmail.com
© 2013 Scott Adams, Inc. /Dist. by Universal Uclick
www.dilbert.com
9-7-13

1986–2016
ESnet
30 YEARS OF NETWORKING
AT THE SPEED OF SCIENCE

Motivations

- My goal is to break down this myth by viewing the Science DMZ *as a security architecture.*
- That is, by thinking about Science DMZ as a form of security *control,* not just something that needs to be controlled.
- At the same time, Science DMZ enables us to do a better job of risk-based security through segmentation.

Risk-based vs. Control-based Security

- Risk-based (ideal form):
  - Identify risks (impact and likelihood over a period of time).
  - Identify and/or create controls that are specifically designed to mitigate those risks.
  - Apply controls as necessary.
- Control-based (ideal form):
  - Select controls from a checklist or standard.
  - Controls are, or at one point were, believed to mitigate a general set of risks.
  - Apply controls (more controls==better security).

Risk-based vs. Control-based Security

- Most security experts prefer risk-based security
  - Control-based security: apply controls "because the standard says so."
  - It's actually hard to find, in the literature, anyone who likes or prefers control based security.
  - Broad application of firewalls (e.g. large border firewall), often viewed as control-based security.
- So why do we still practice control-based security in many instances?
  - Risk based security is actually pretty hard.
  - Risk assessment itself is hard.
  - Determining if a risk is actually being mitigated is hard.

Risk-based vs. Control-based Security

- The non-falsifiability of security assessments (Microsoft Research paper):
  - Indicates difficulty with fully assessing risk (but also effectively dismisses control-based security).
  - In simple terms, it's easy to find cases where a security breach *wouldn't* have happened if a particular security control were in place, but it's pretty much impossible to say that a security breach that didn't happen, would have happened, if a security control hadn't been in place.
  - Early days of firewall logging: "Our firewall prevented 1,789,034 attacks last week!"

Risk-based vs. Control-based Security

- Other things that make risk-based security hard:
  - It’s labor-intensive.
  - It may be more expensive up-front, but likely cheaper in the long run.
  - Rumsfeld’s razor: What about all of the unknown unknowns?
  - “Nobody ever got fired for having a firewall.”
- Moreover: **The set of risks at a research lab or university campus demonstrably vary across the resources that are attached to the network.**
- However, this turns out to be more of an argument against control-based security.

Network Segmentation

- Think about your residence hall networks, business application networks, and the networks that are primarily in research areas.
- The risk profiles are clearly different, so it makes sense to segment along these lines.
- Your institution may already be doing this for things like HIPAA and PCI-DSS. Why? *Because of the controls!*
- The Science DMZ follows the same concept, from a security perspective.
- An example here is how using a Science DMZ to segment research traffic (especially traffic from specialized research instruments) can actually *improve* campus security posture.

Network Segmentation and the Science DMZ: An Example

- I typically look at two examples:
  - Scenario 1: Scientific Instruments
  - Scenario 2: HPC clusters

Scenario 1: Scientific Instruments

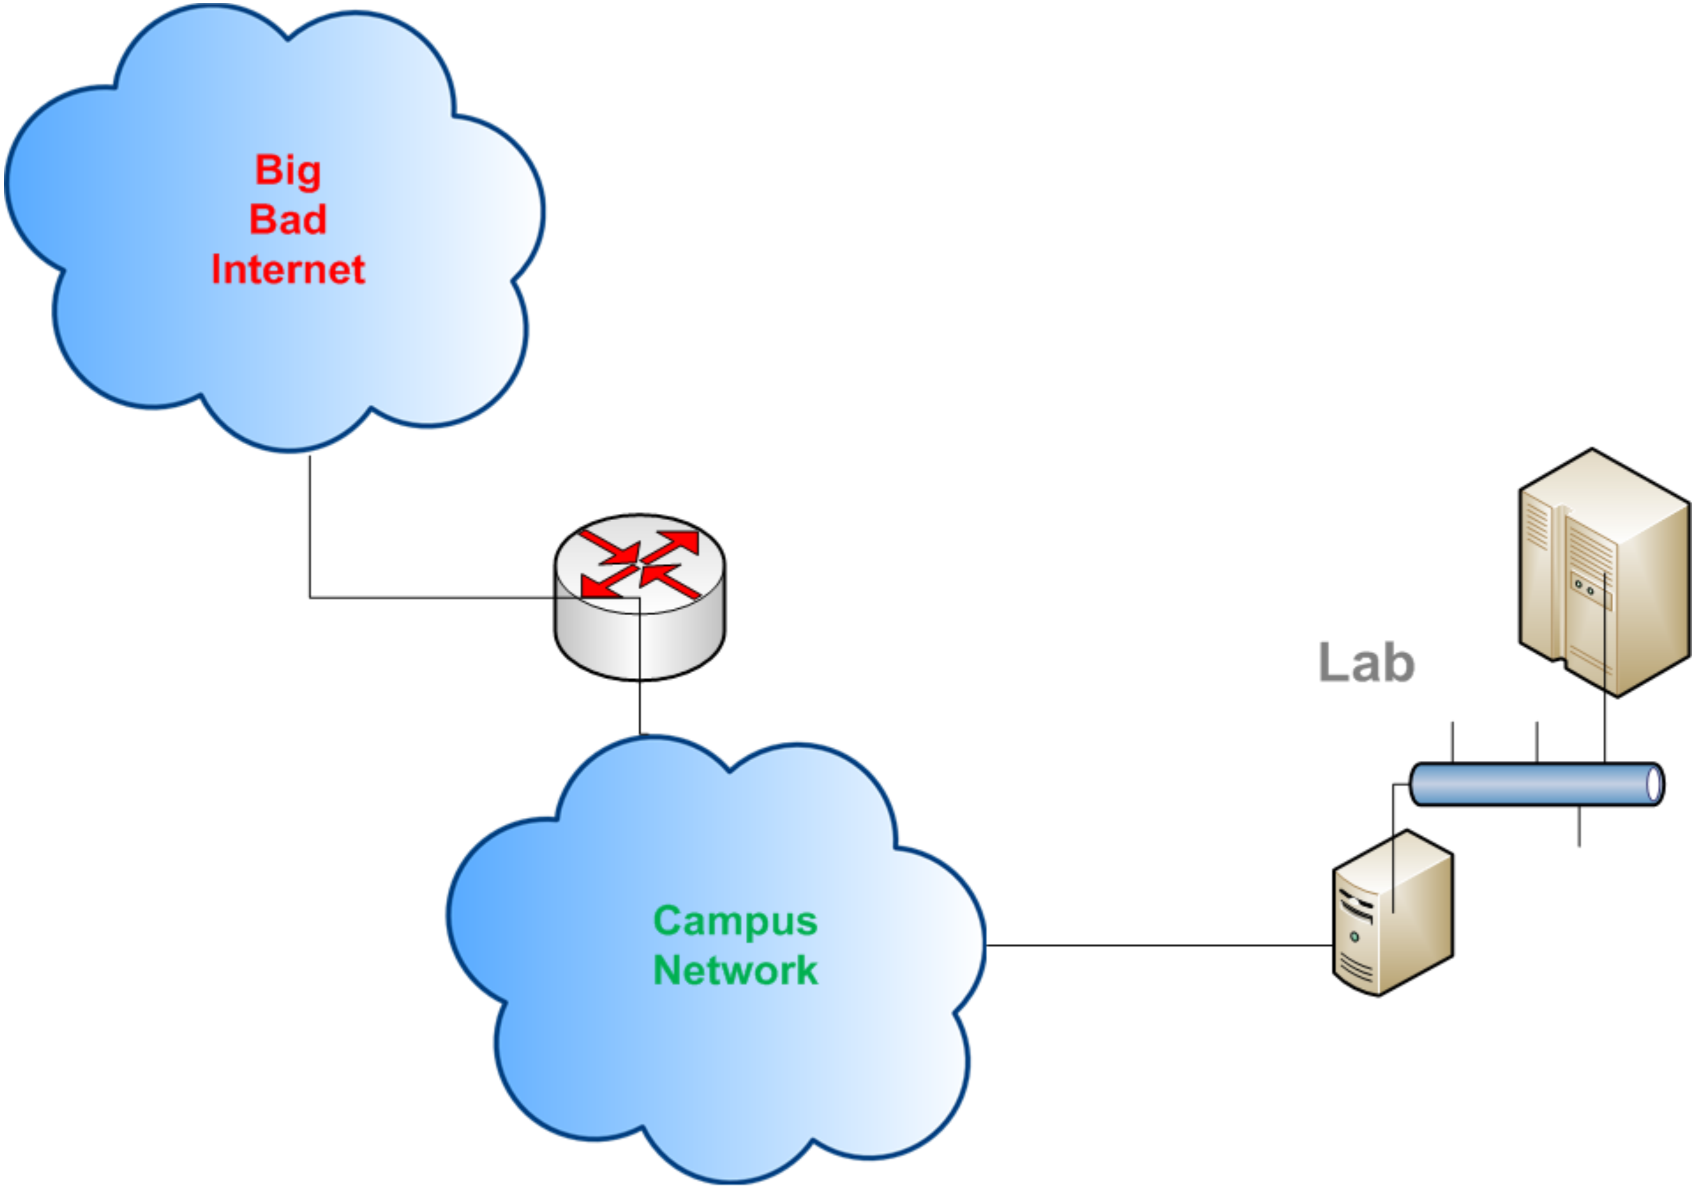

Scenario 1: Scientific Instruments

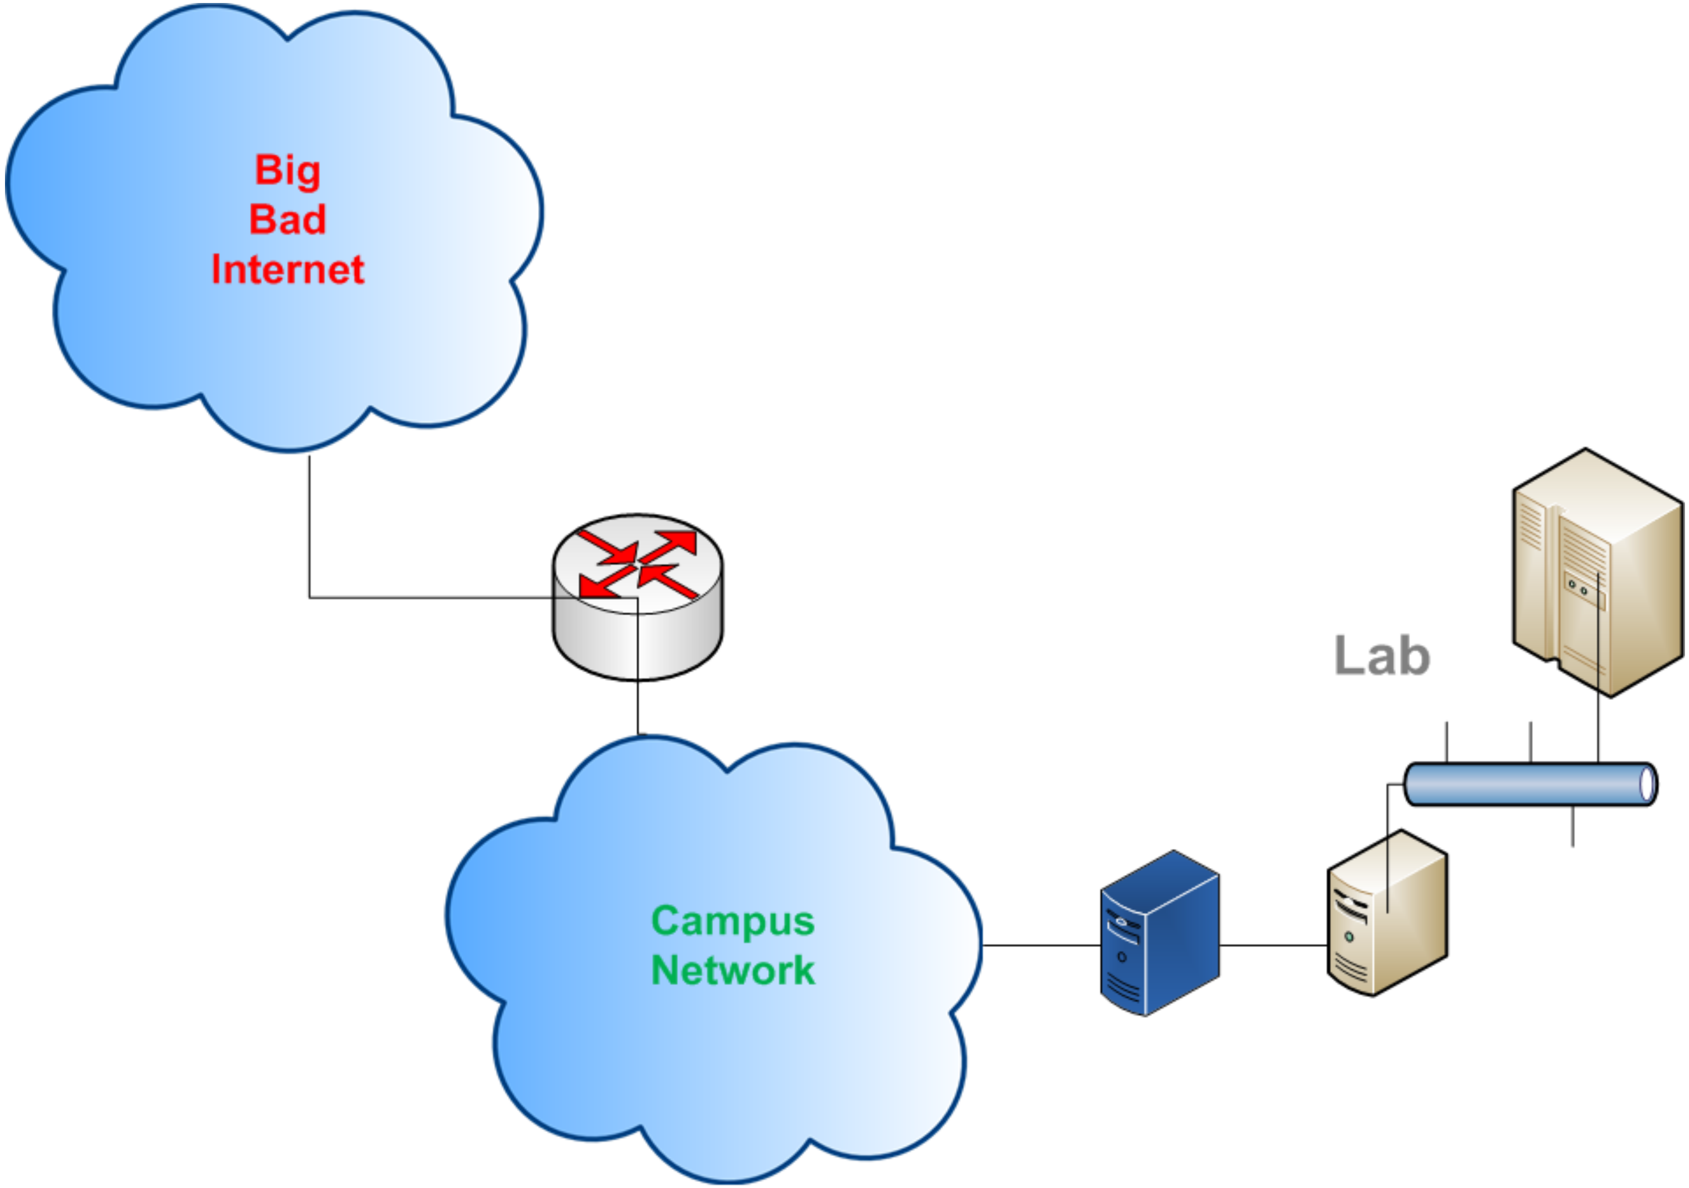

Scenario 1: Scientific Instruments

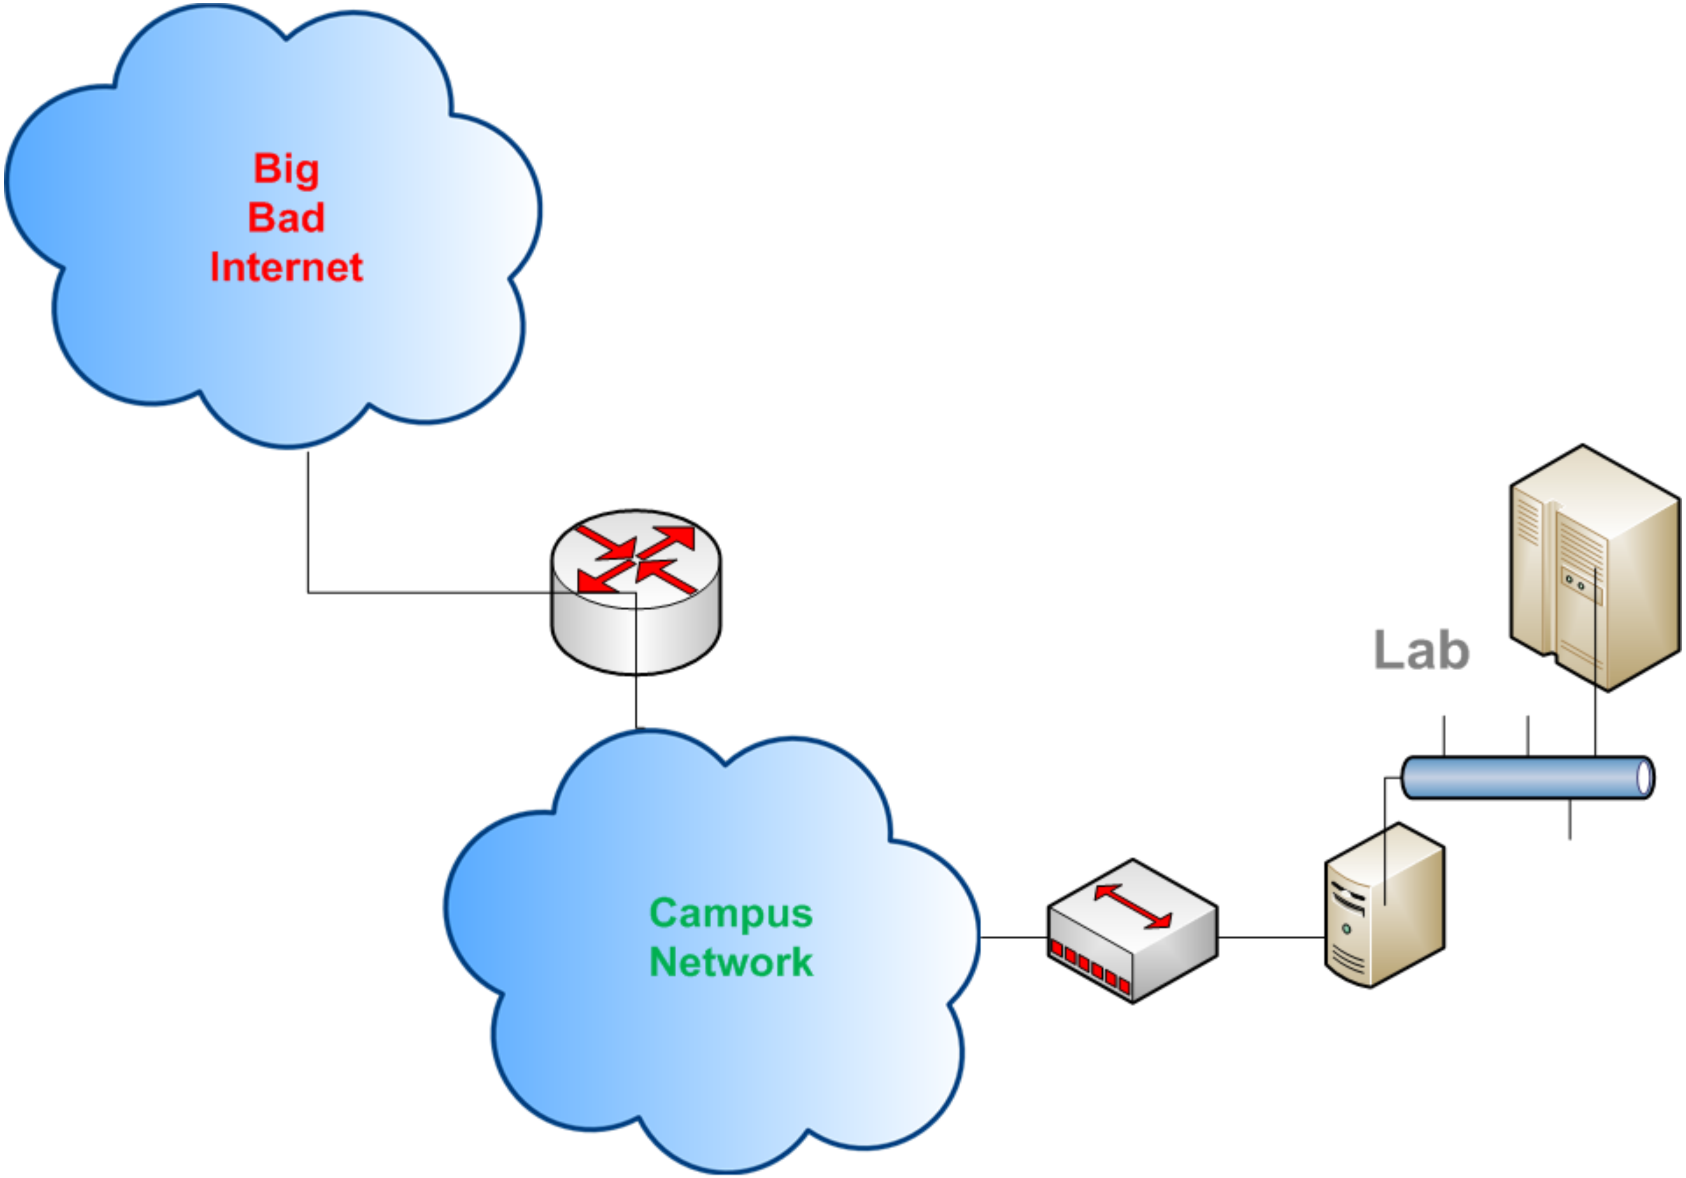

Scenario 1: Scientific Instruments

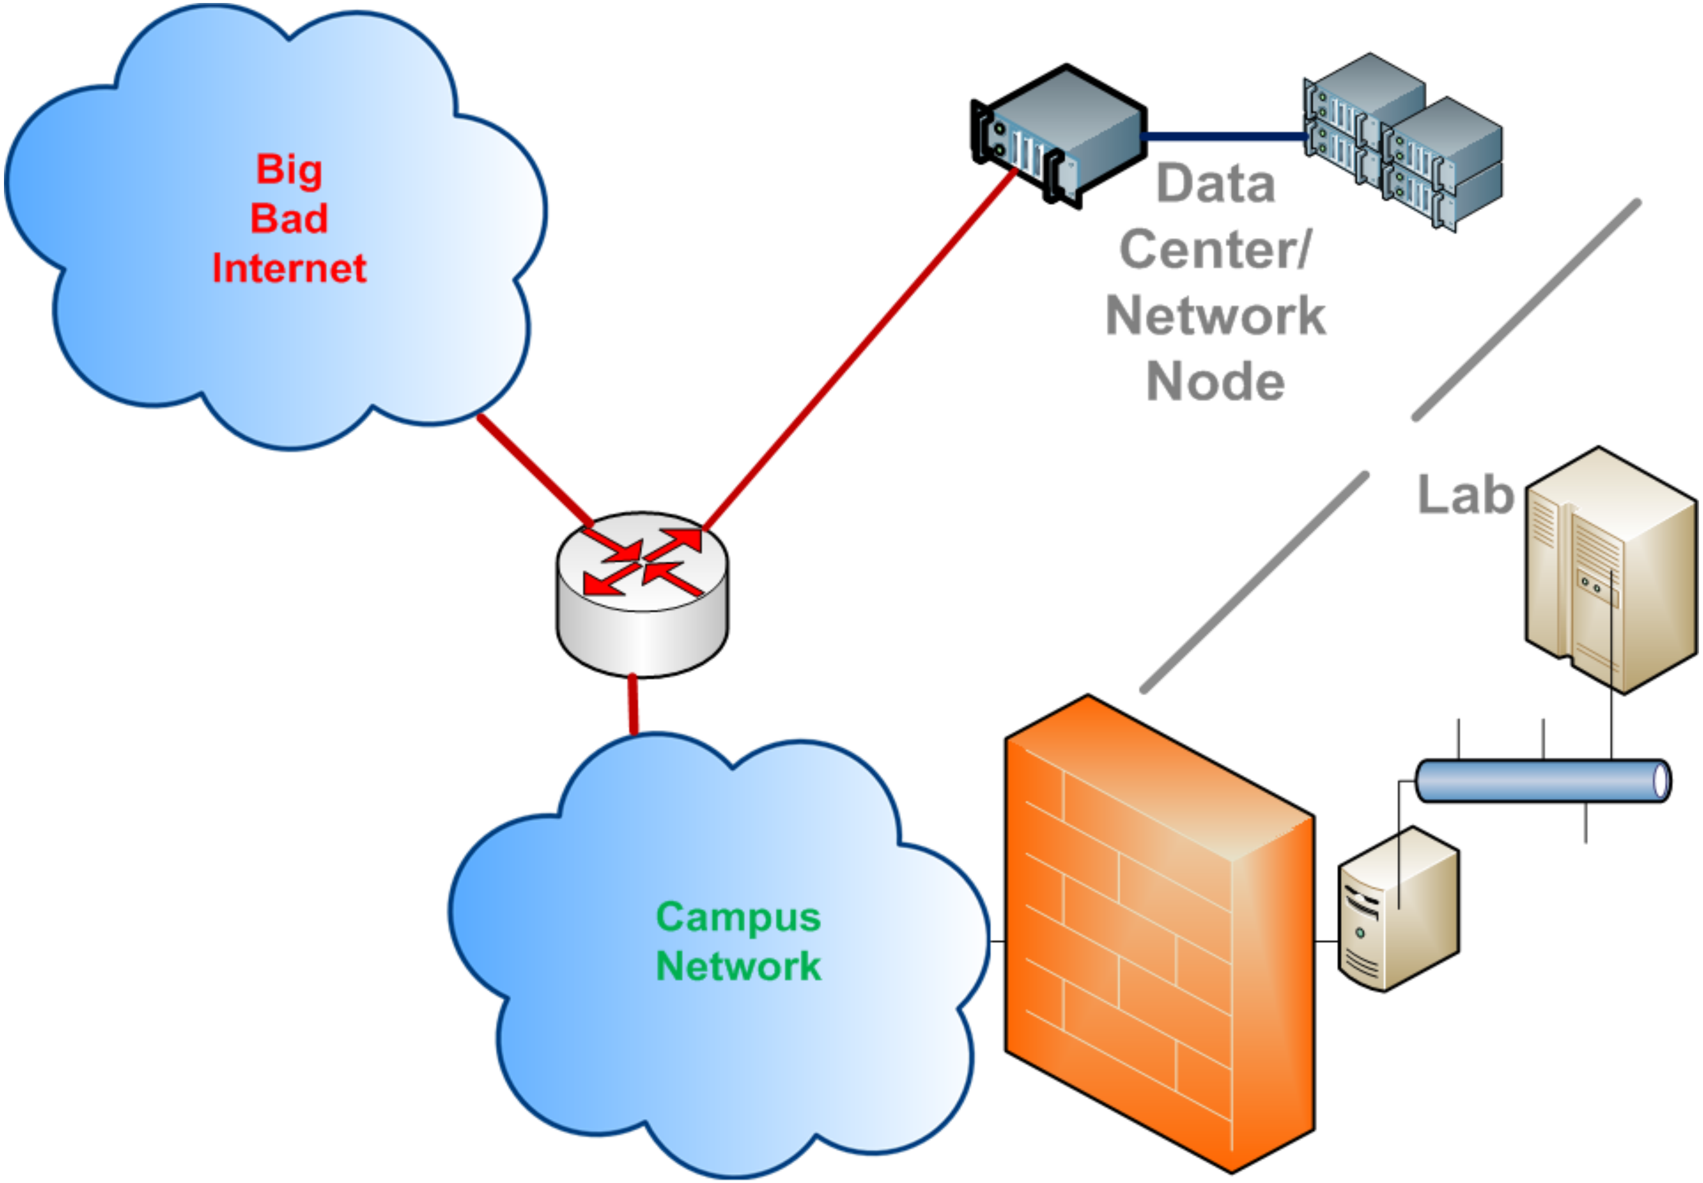

Scenario 1: Scientific Instruments

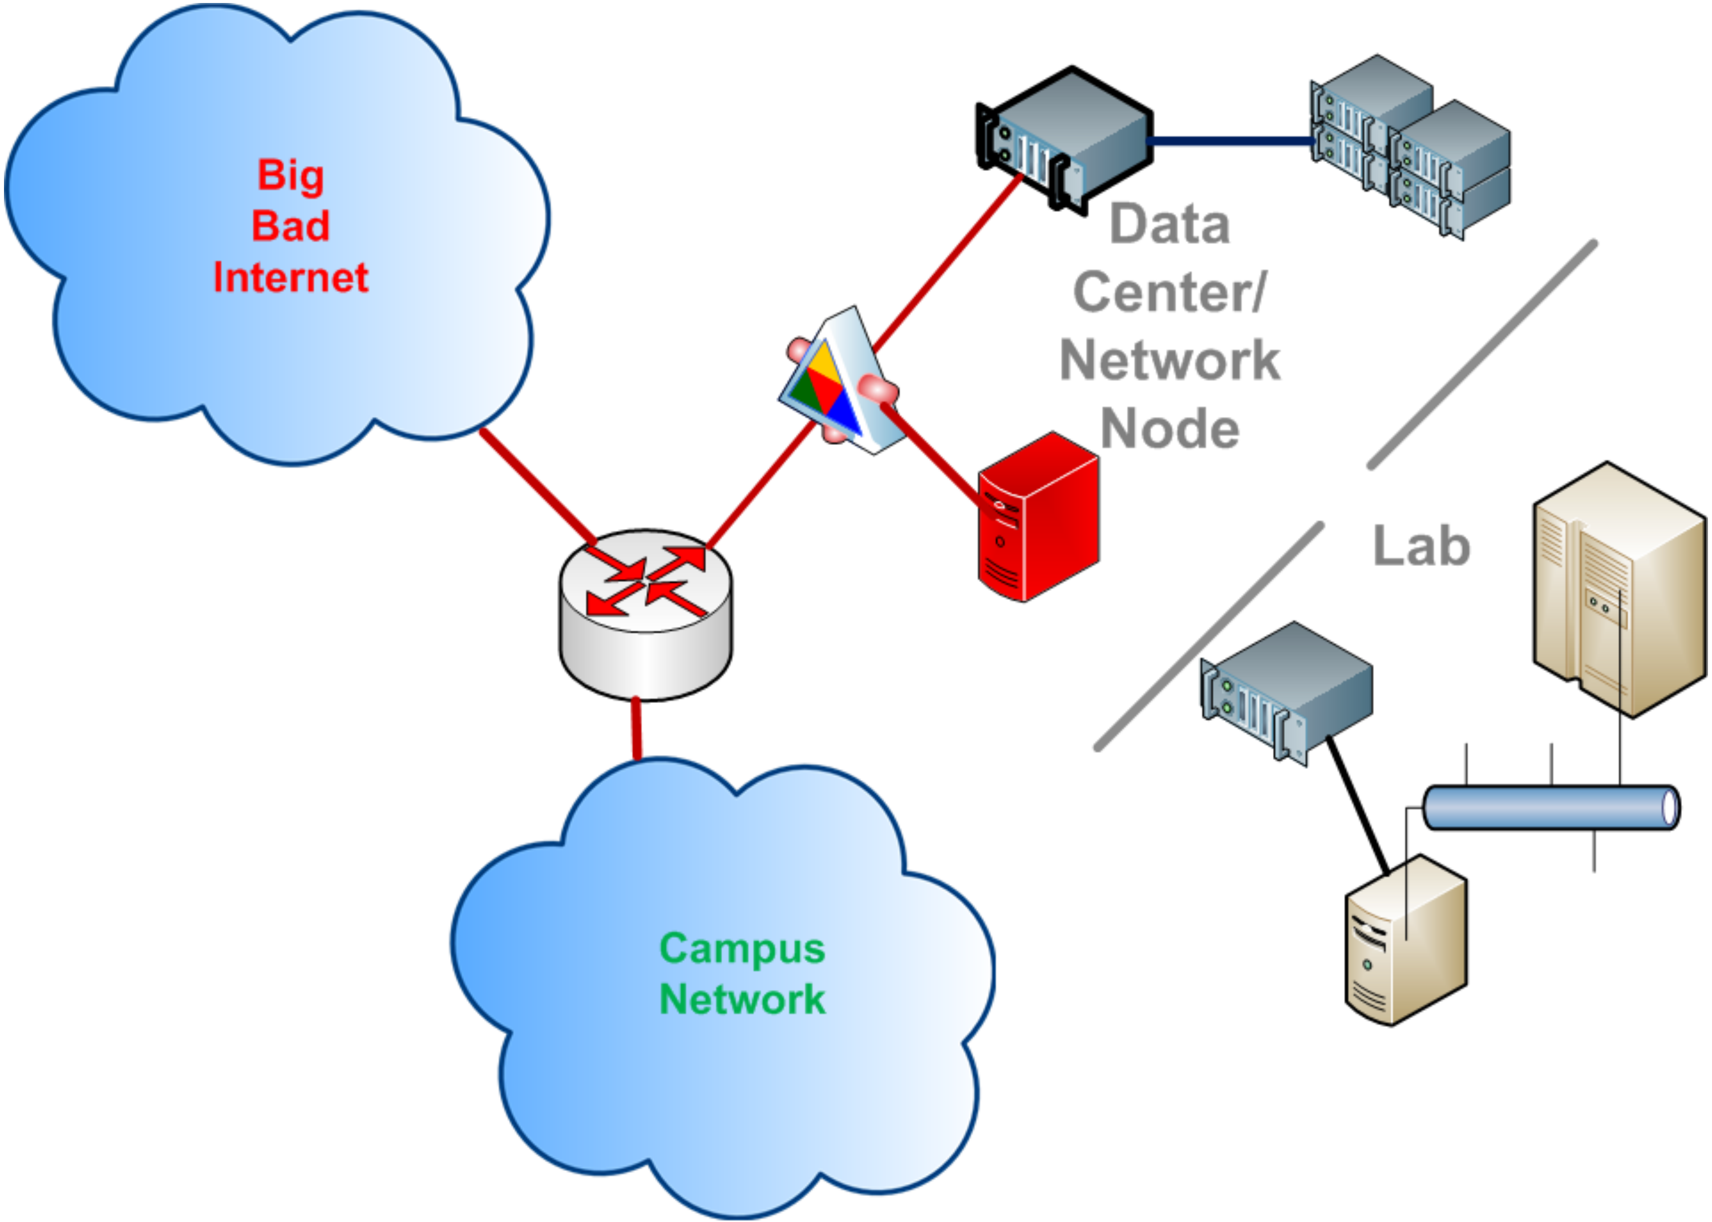

Scenario 1: Scientific Instruments

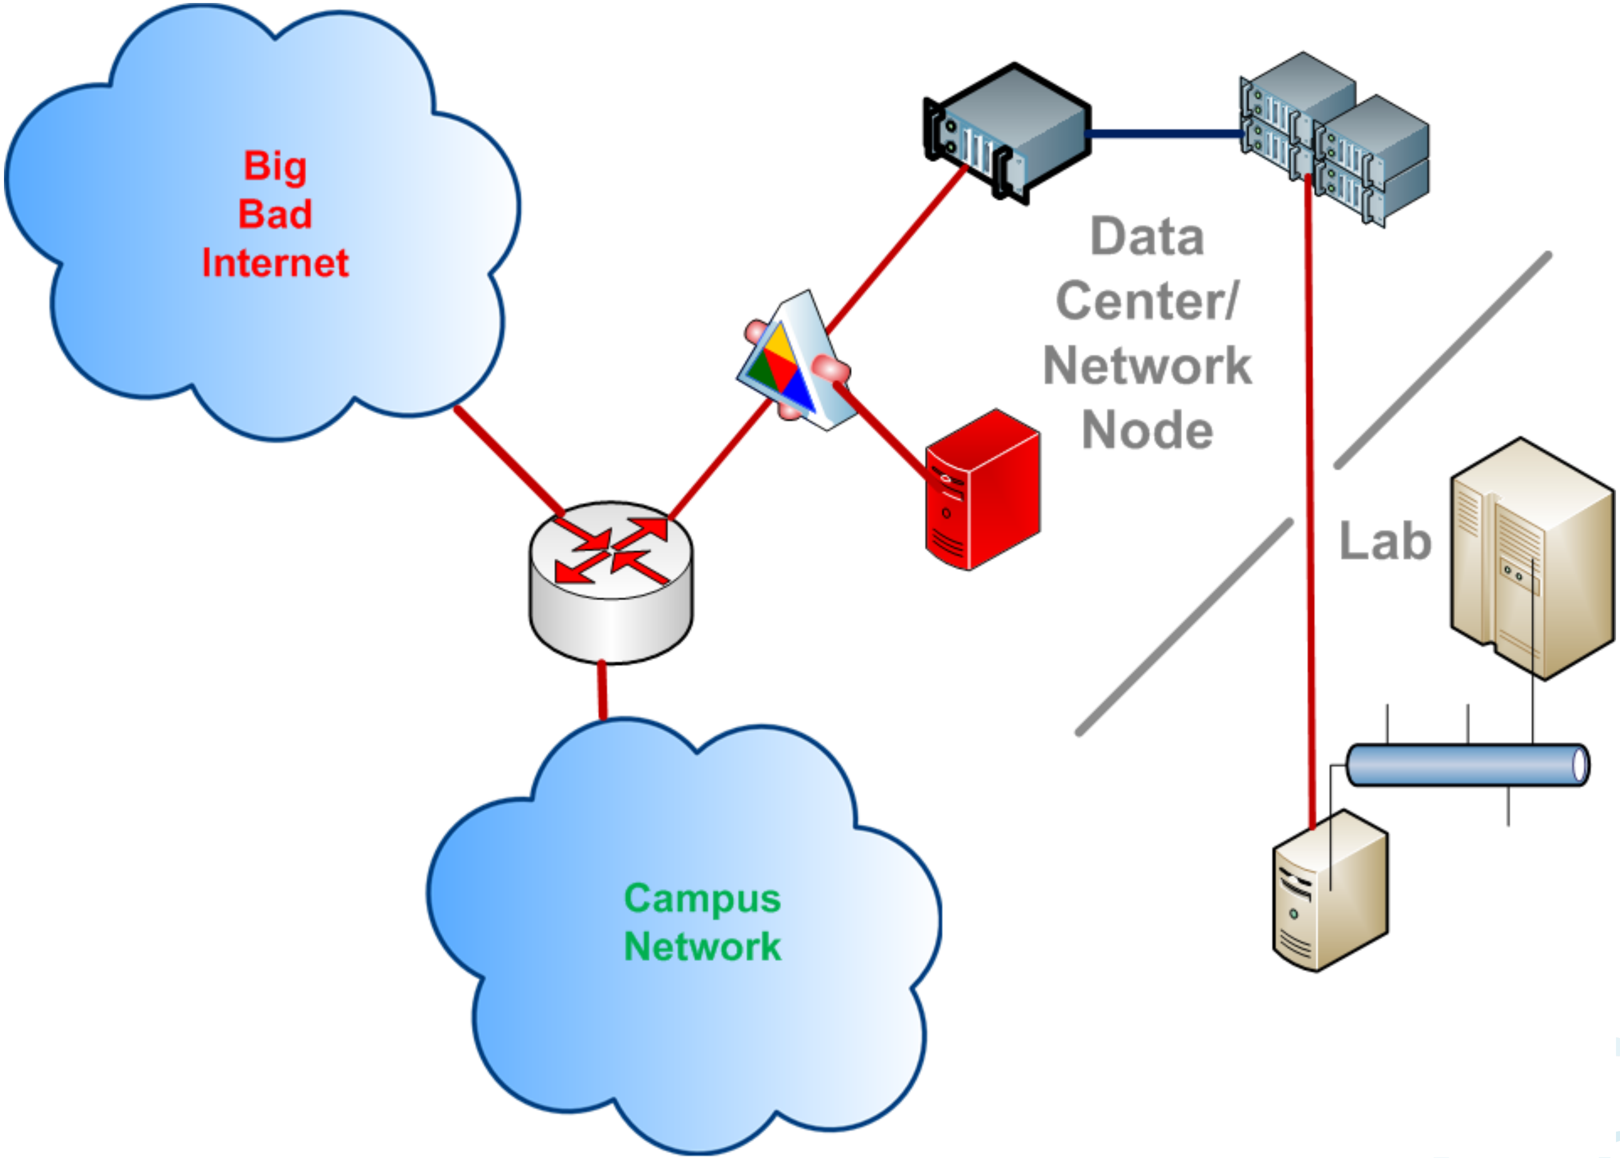

Scenario 1: Scientific Instruments

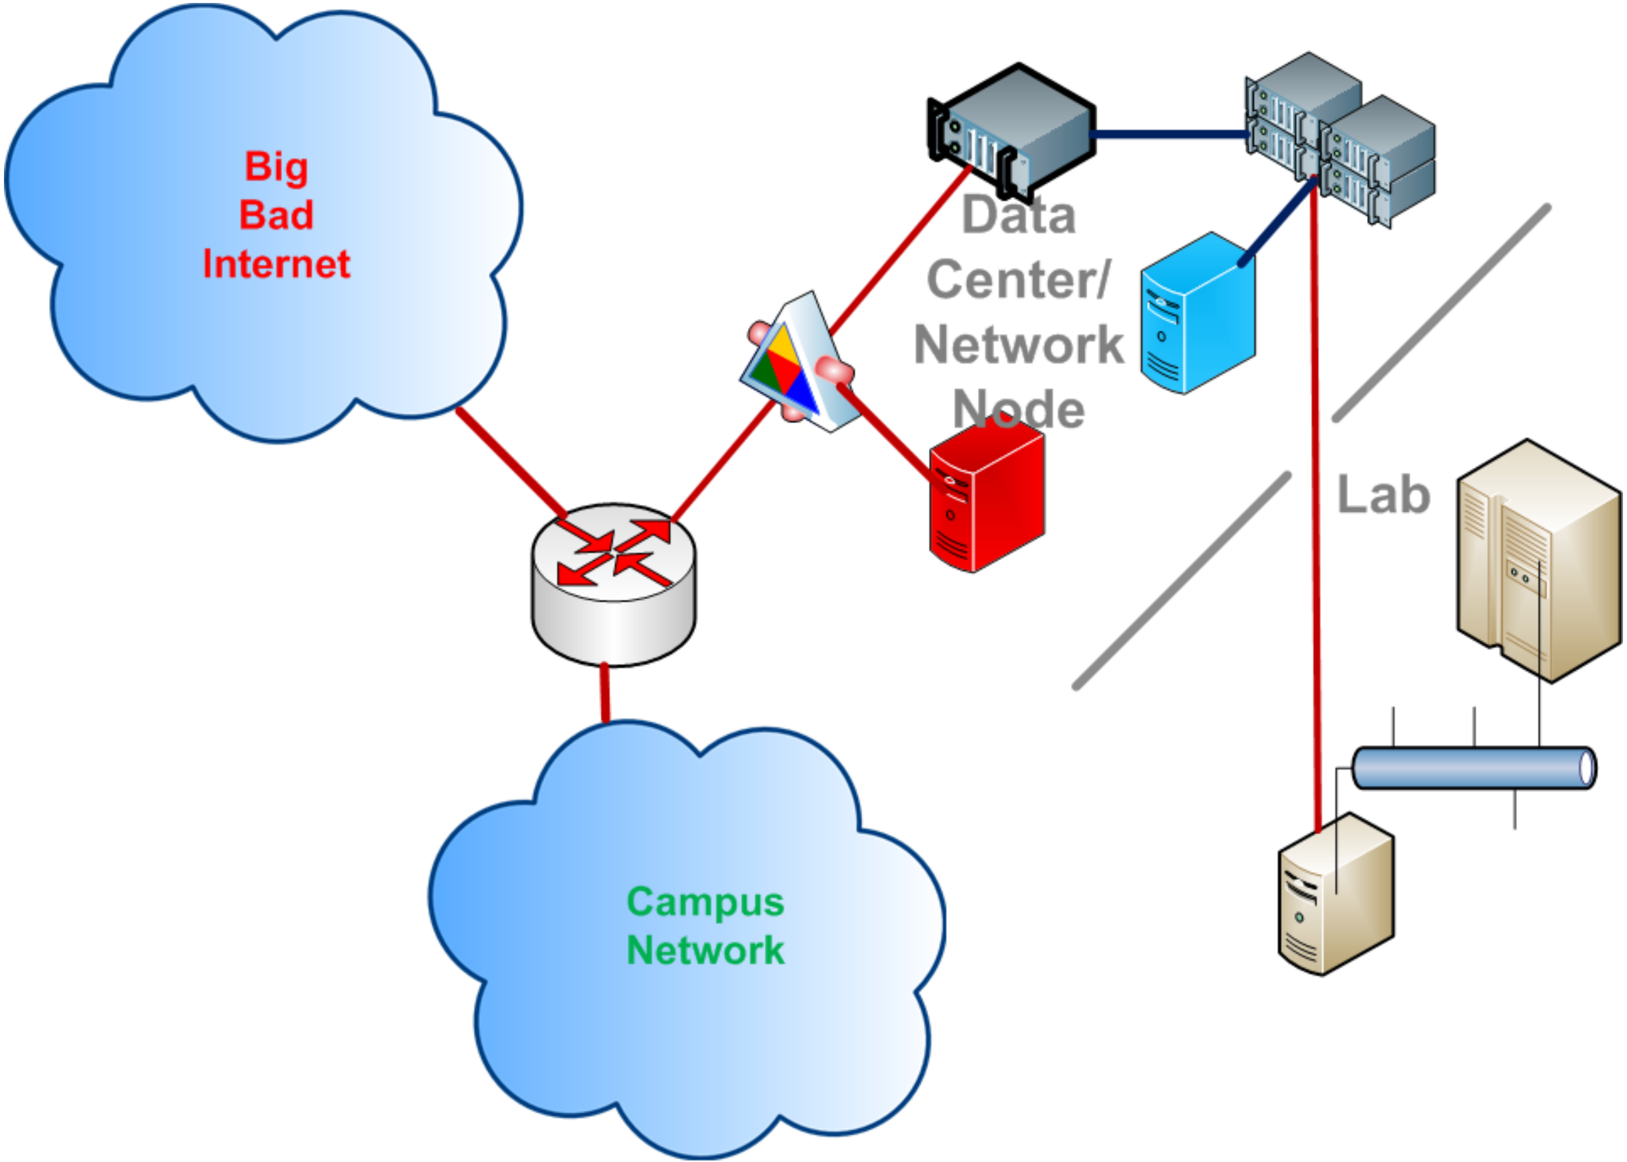

Scenario 1: Scientific Instruments

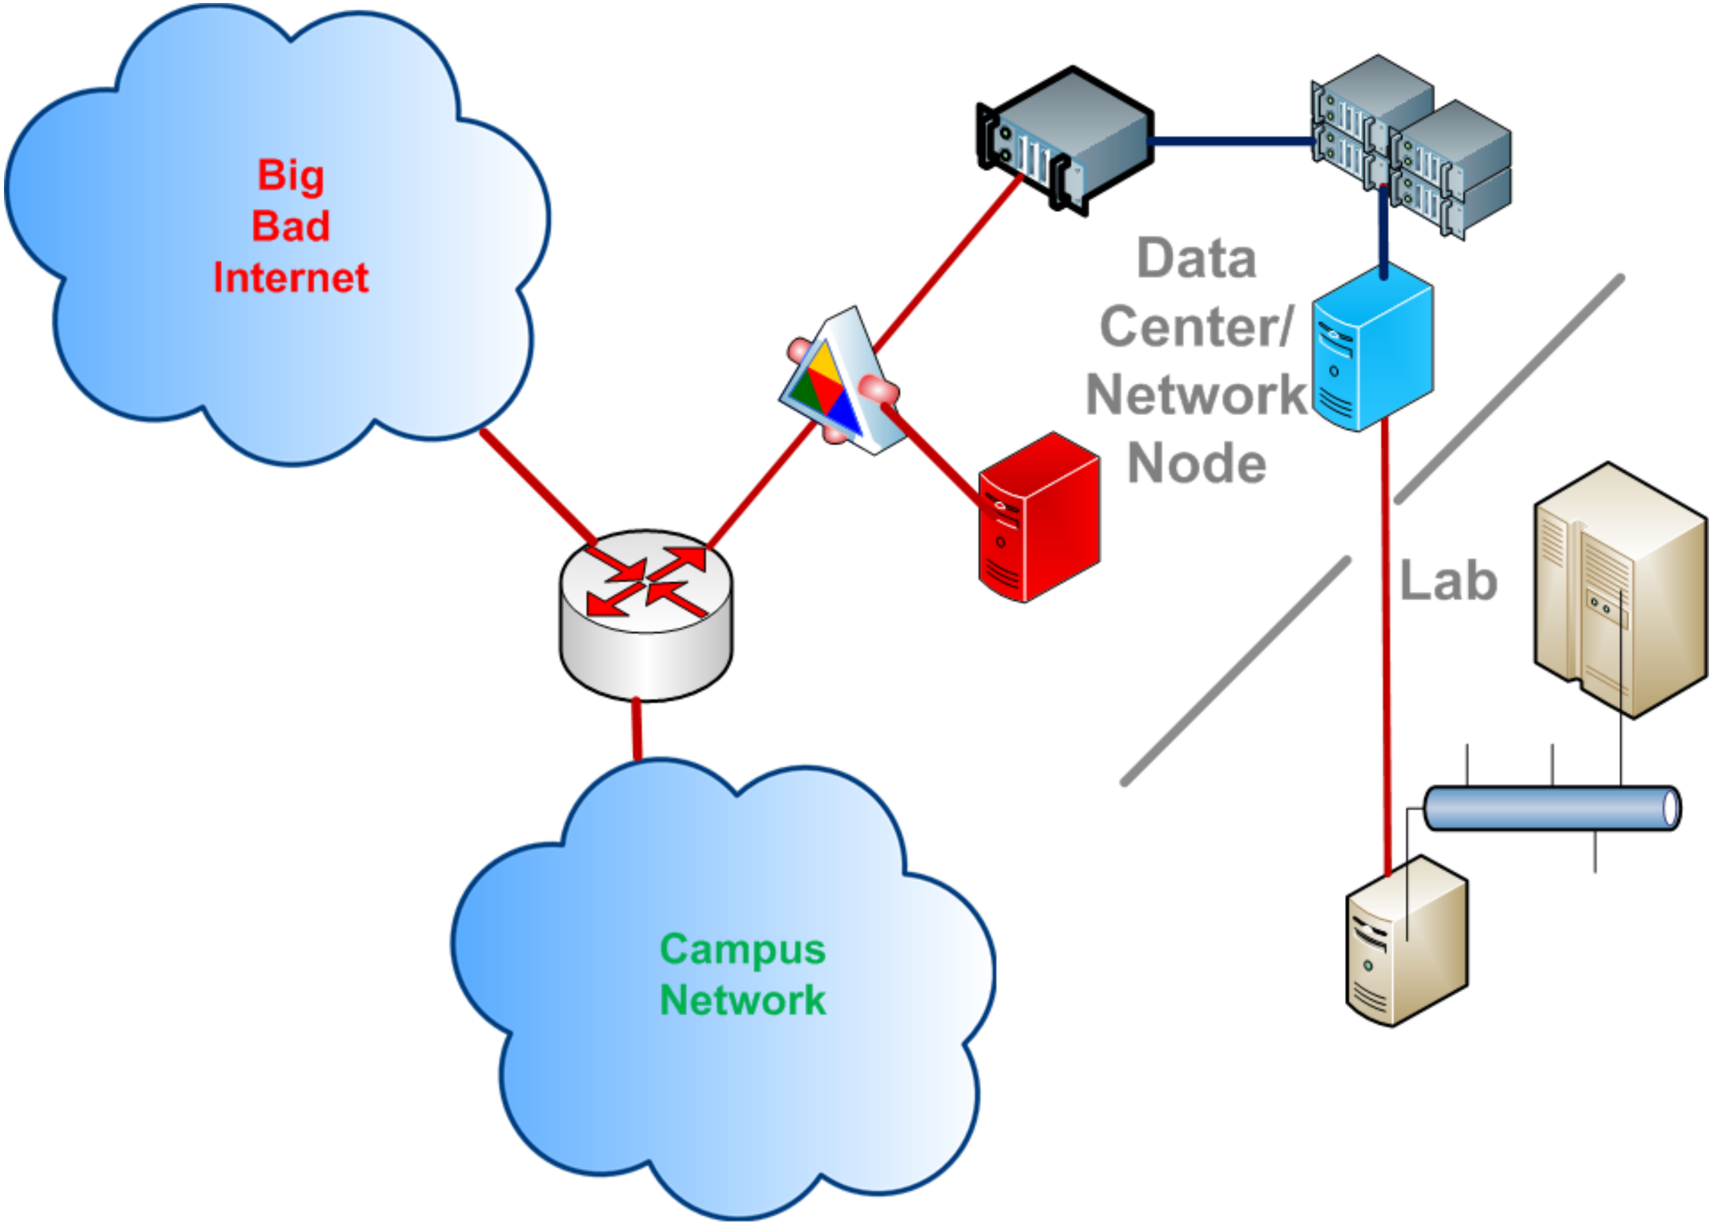

Scenario 1: Scientific Instruments

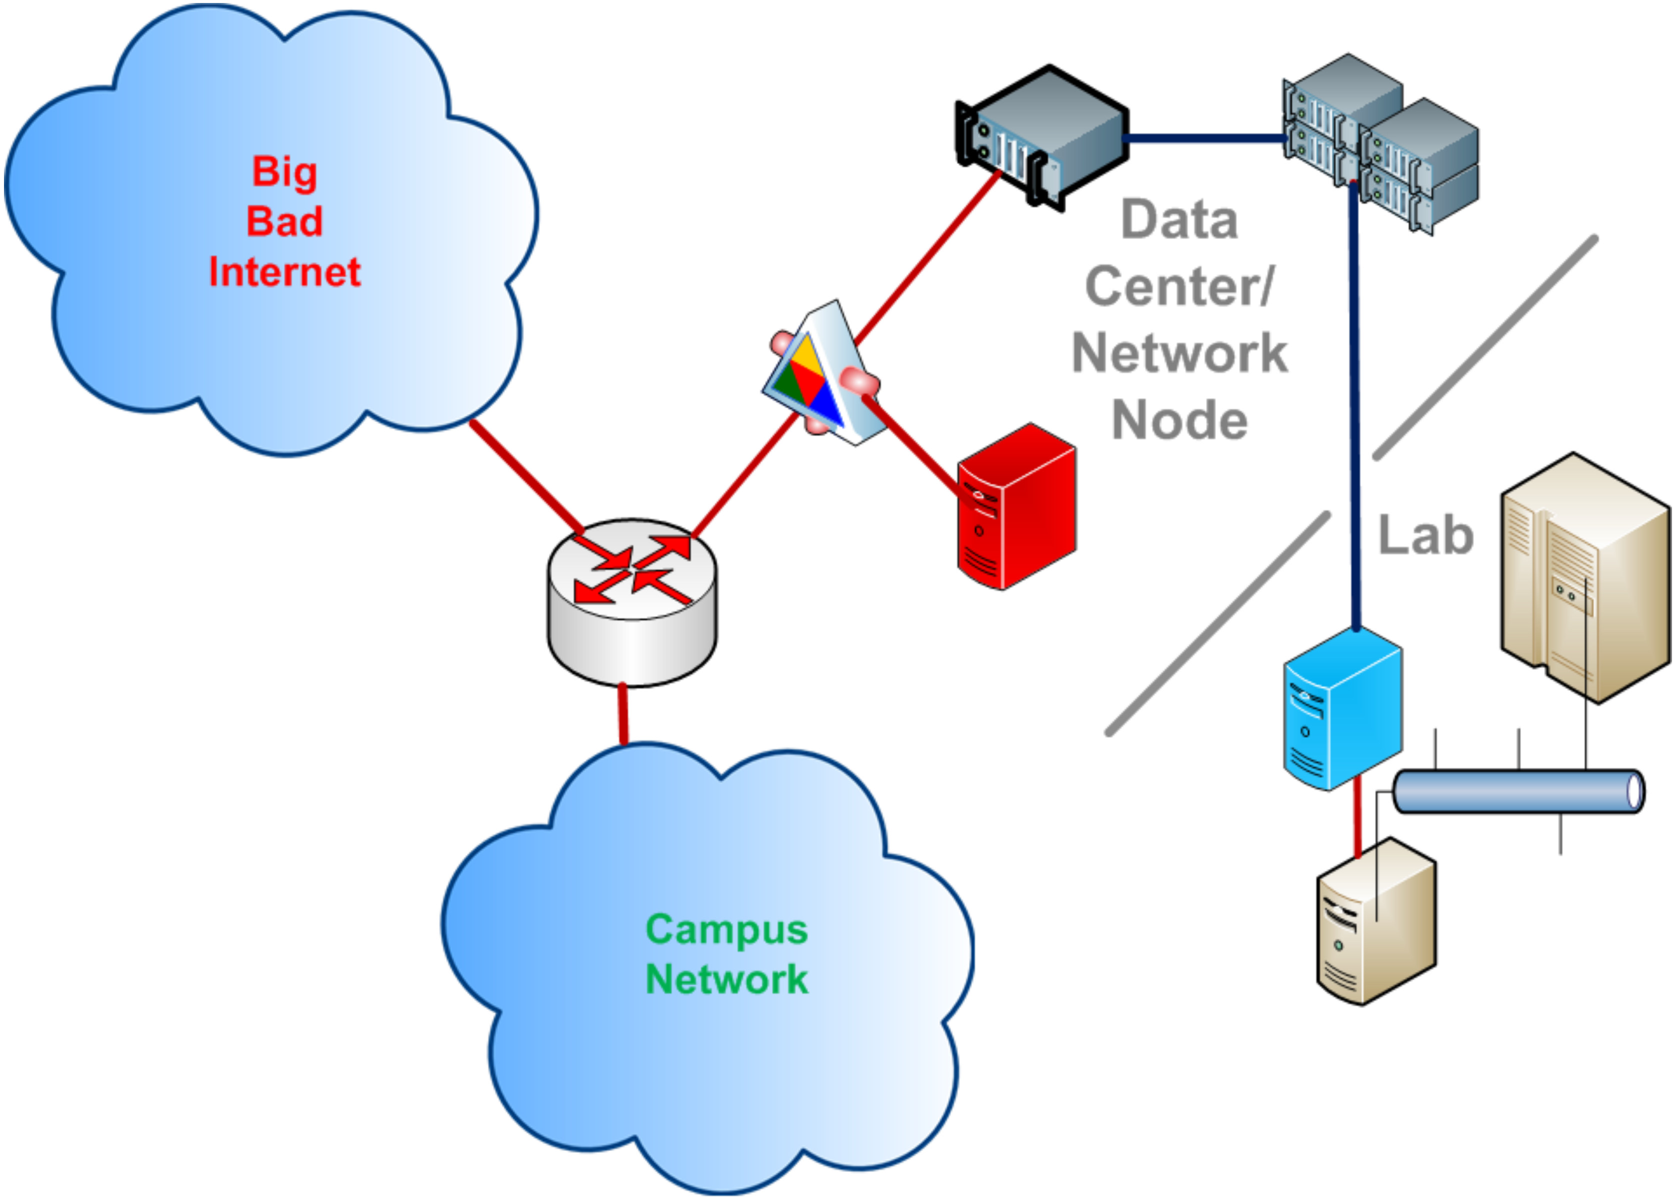

Scenario 2: HPC Clusters

- 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: HPC 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 and Implications

- Think about what the Science DMZ is trying to do.
  - Improve performance, both by removing impediments and improving the performance of the devices that must be in line.
  - Ease troubleshooting.
  - In general, reduce degrees of freedom from science networks.
  - Maximize performance **and** security **and** resiliency.
- A lot of campuses are building ”distributed Science DMZs” or “Science Networks.” These are good, but they may not realize the full benefit.
- When I think about the problems we are trying to solve, I still wonder if layering “SDN” on top will be an answer (let alone “the” answer).