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

Internet2 Joint Techs

28 September 2016

U.S. DEPARTMENT OF ENERGY
Office of Science

BERKELEY LAB

Firewalls: To Control or Be Controlled?

Motivations

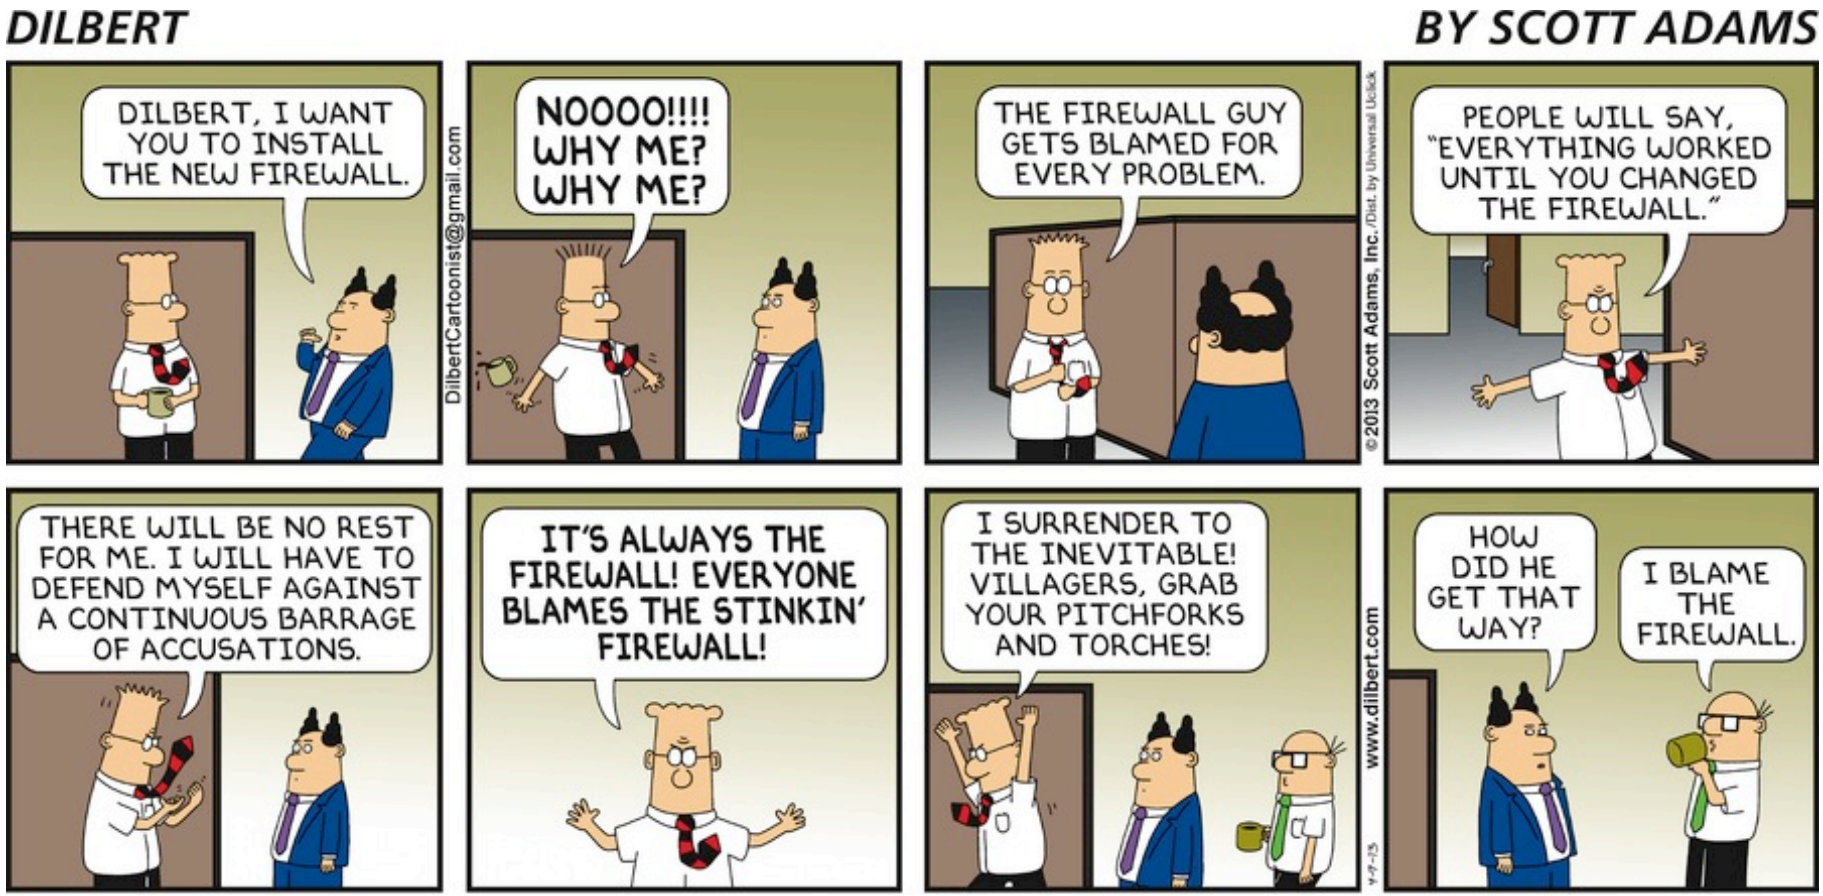

Motivations

- 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.

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.

Examples and Scenarios

- See the longer talks!
  - Science DMZ Security: Presented to CENIC 2015: http://registration.cenic.org/cenic2015slides/ScienceDMZSecurity-Buraglio--Sinatra-dart-v2.pptx
  - CTSC Presentation: http://trustedci.org/s/trustedci-sciencedmz-security.pdf and nice long webinar archive (https://ctsc.adobeconnect.com/dmz_recorded/event/registration.html)
- CTSC Cyberinfrastructure Security Workshop paper forthcoming as part of proceedings.

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).

Citations and Acknowledgements

- Reviewers: Jason Zurawski, Eli Dart, Mike Dopheide, Denise Sumikawa, Sam Oehlert, Nick Buraglio, Lauren Rotman
- Conferences and workshops: CENIC, CTSC, Von Welch, Jeanette Dopheide, Jim Basney
- Science DMZ Security: Presented to CENIC 2015: http://registration.cenic.org/cenic2015slides/ScienceDMZSecurity-Buraglio--Sinatra-dart-v2.pptx
- CTSC Presentation: http://trustedci.org/s/trustedci-sciencedmz-security.pdf and nice long webinar archive (https://ctsc.adobeconnect.com/dmz_recorded/event/registration.html)