

Science DMZ Security Architecture

Michael Sinatra (presenting)

Eli Dart (data portal content)

Nick Buraglio (Bro content)

Network, Outreach and, Security
Engineers

Energy Sciences Network

Lawrence Berkeley National

Laboratory

ASTRON Workshop

Dwingeloo, NL

24 September 2018

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...



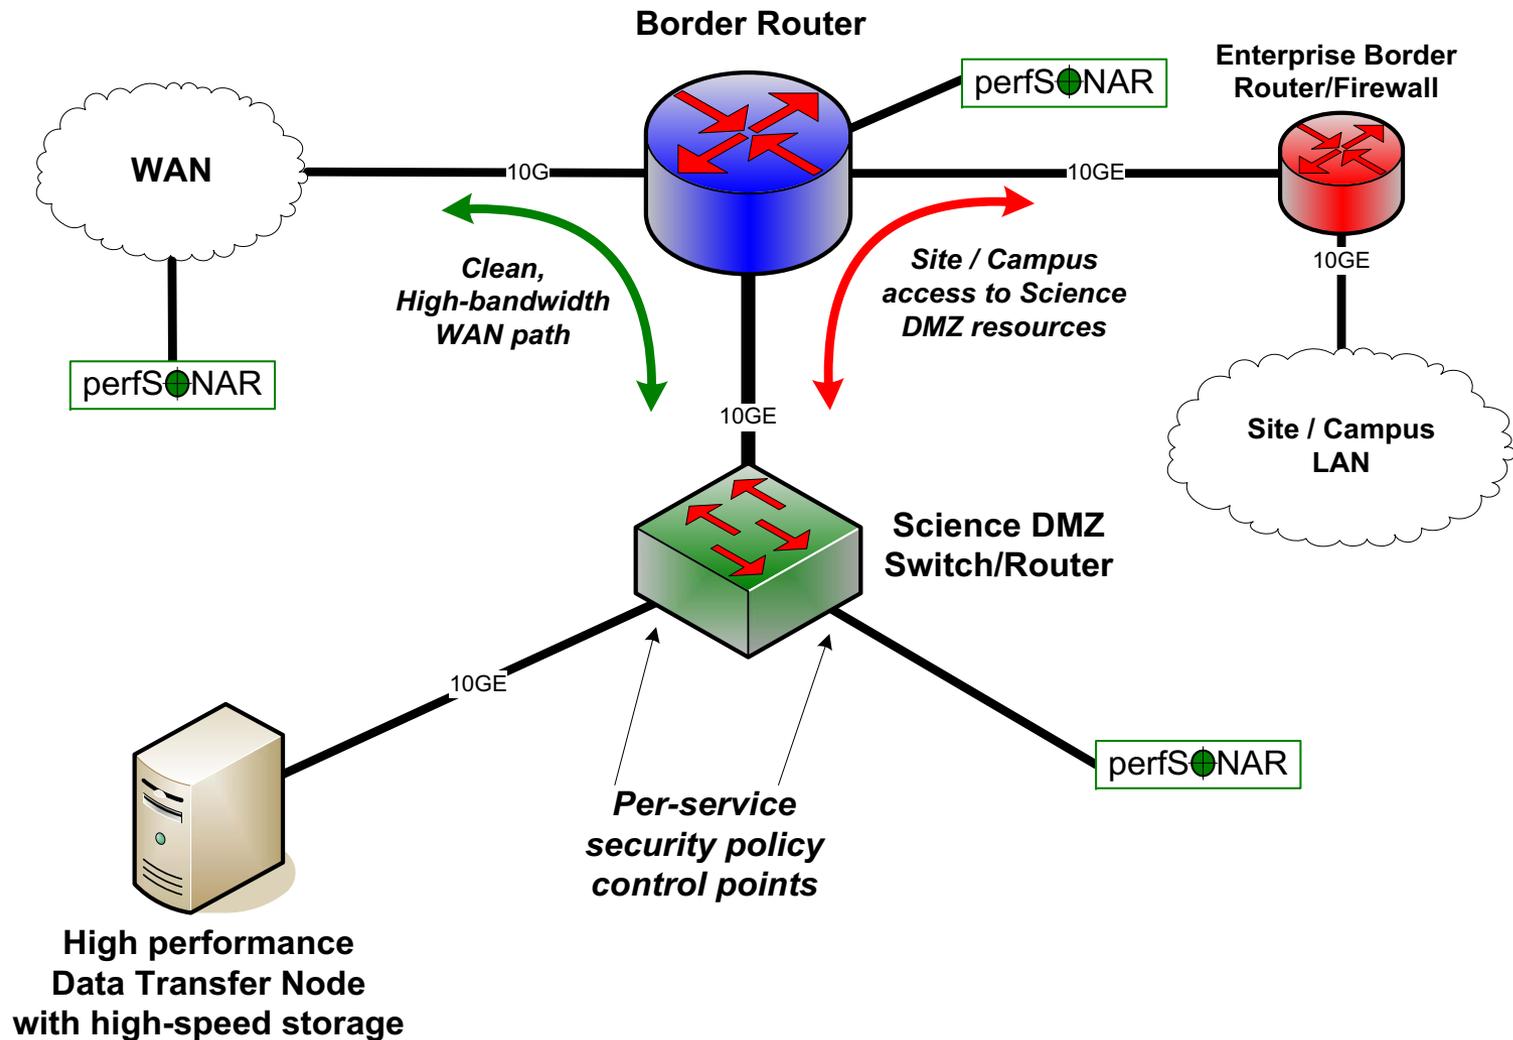
Who am I? Why am I here?

- ESnet is operated at Lawrence Berkeley National Lab for the US Department of Energy.
- Serves the entire US National Lab Complex, plus one of the service providers that connects to CERN.
- DOE one of the largest public funders of basic science in US.



Classic Science DMZ Design

Science DMZ Design Pattern (Abstract)



How Network Engineers and Scientists View the Science DMZ



How “Security People” View the Science DMZ



What the Science DMZ Really Is



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

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

- Segmentation also allows more granular *trust* between services.
- The Science DMZ does not have to be trusted by the rest of the campus/laboratory network.
- Many US EDUs and labs implement the zero-trust model for Science DMZs.
- Science DMZ is treated as “outside the perimeter” for most campus services.

Securing the Science DMZ – Contain risk

- The Science DMZ should run a limited number of services so that the risk of each service is easily contained, controlled, and mitigated.
- In general, the DTN is the primary function of the Science DMZ.
- Science DMZs should not have email servers, web servers, blogs, forums, XMPP or other chat servers (with the possible exception of ChatOps-like functions), and other non-DTN stuff.
- Keep the application software simple, so that it's easier to maintain.
- For applications like data portals, there is a potentially better solution--discussed later.

Default-deny? Yes!

- Remember, the Science DMZ model is consistent with a default-deny access policy.
- The only difference is that this policy should be statelessly applied at the router.
- Stateful inspection and (especially) deep-packet inspection can impact performance, but by constraining the function of the Science DMZ, stateless rules can provide a major gain in protection.
- Some firewalls will convert stateful rules into stateless dynamic rules and can transfer data at line rate.
- Others have SDN-like functionality to route around the packet inspection engine for “science flows.” But beware...

Complexity is the enemy

- Layering SDN functions on top of a Science DMZ increases complexity. Complexity is bad on several levels...
 - Can directly lead to operational (and therefore security) problems.
 - Makes troubleshooting harder. Remember, one major point of the Science DMZ is to reduce degrees-of-freedom and make troubleshooting easier for data transfer applications.
 - Heuristics still leave a lot to be desired. How do we identify “science flows” properly and reliably? How do we do it without trampling on them first?

Other things you can do...

- Implement black-hole routing and/or BGP Flowspec (RFC 5575) and tie into intrusion detection system. BGP flowspec rules should be able to provide line-rate protection (just like ACLs), and subject to the same limitations.
- Host-based firewalls...
 - Yes, they do work.
 - Yes, they do perform, although Linux users may want to look into nftables as a (much-needed) replacement for iptables.
- Implement IPv6 on your Science DMZ.

Other things you can do...

- Limit routing to/from the Science DMZ. For example, ESnet's DTNs only route to known R&E entities and AWS. (See, for example http://stats.es.net/sample_configs/pscheduler/).
- Outbound ACLs (similar to limiting routing).
- Automation and central management of...
 - Account management
 - Host-based firewalls
 - Auto patching of software
- *Intrusion detection and monitoring*

What if you have sensitive or restricted data?

- Encryption is key
 - Most modern transfer tools (e.g. Globus tools) will encrypt data in flight.
 - I would (also) recommend encryption at rest—that is the best way to protect the *valuable thing*.
- Data should be encrypted before being accessed by the Science DMZ.
 - Put data on shared filesystem already encrypted – could be encrypted by a dedicated system;
 - or, save data to unshared filesystem, encrypt, unmount LUN, and then have the Science DMZ mount it.

What if you have sensitive or restricted data?

- Examples of medical science DMZs:
<https://academic.oup.com/jamia/article/25/3/267/4367749>
- We'll also talk later about the use of *data portals*, which can further separate the functions within the data workflow.
- This concept is used in at least some medical science DMZs.

Intrusion detection

- Two major components in most R&E entities:
 - Bro:
 - Packet processing engine and event handler
 - Works as an IDS, but different from signature-based IDSes
 - Highly extensible policy language
 - Can basically be taught to handle many kinds of events, not just security events
 - Signature-based IDS:
 - Snort
 - Suricata
- Yes, you can run both on your campus and in your Science DMZ!

What is the Bro IDS?

- An actively developed intrusion detection system originally developed and published by Vern Paxson in 1998, with work starting as early as 1995 currently funded by the NSF and supported by joint efforts at the International Computer Science Institute (ICSI) and National Center for Supercomputing Applications (NCSA)
- Open Source Software, licensed under the BSD license.
- <http://www.bro.org/>

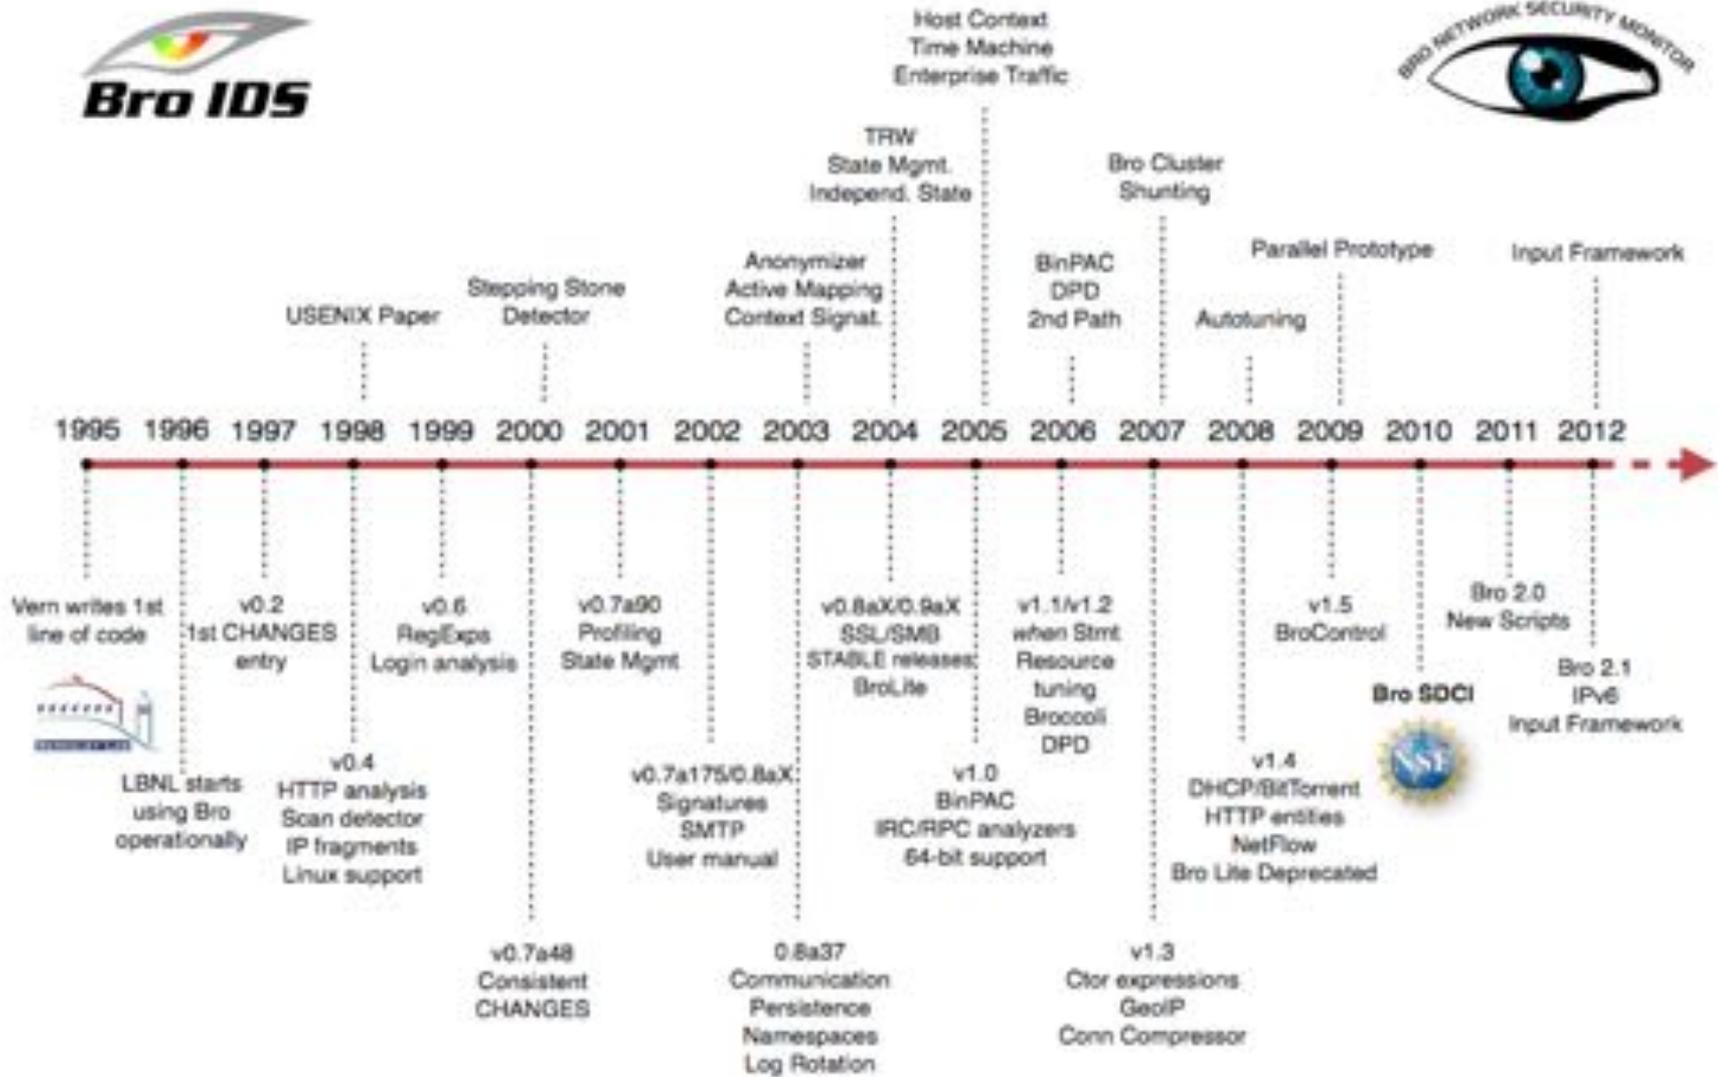


Image source: <http://www.bro.org>

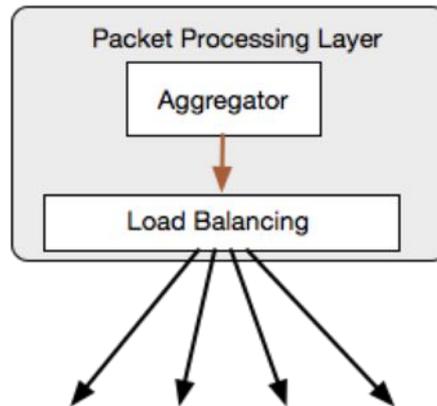


What is the Bro IDS?

- A network Monitoring Platform
 - Commonly used as a power anomaly and intrusion detection system (IDS)
- A modular software stack: three components
 - Packet processing layer
 - Event Engine
 - A policy script interpreter

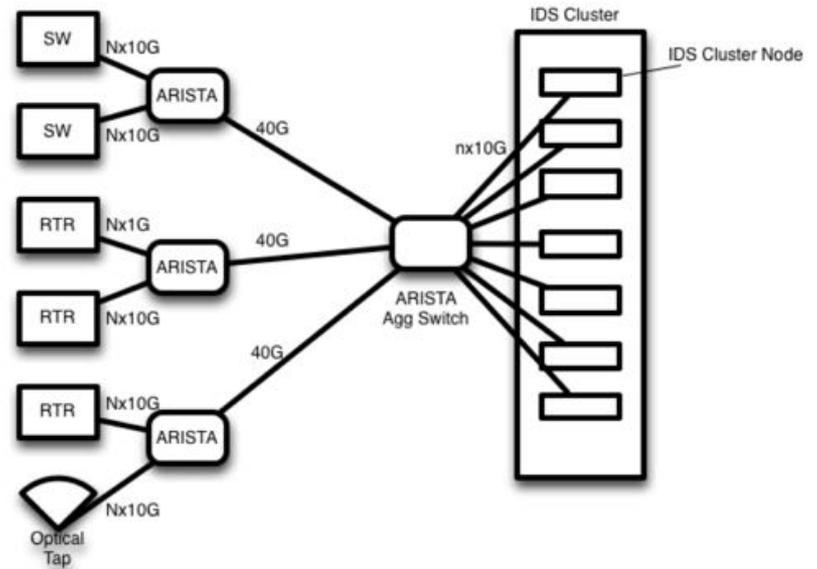
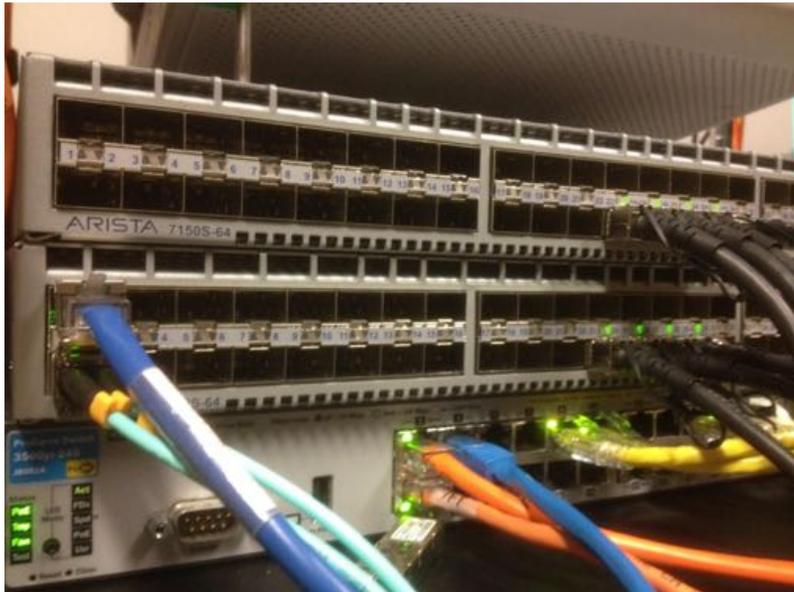
What is the Bro IDS?

- Packet processing layer
 - Has knowledge of what the higher layers need
 - Can exist as hardware or software
 - Pass data to higher layers according to configuration / policy
 - In most cases this layer is an external device or software stack



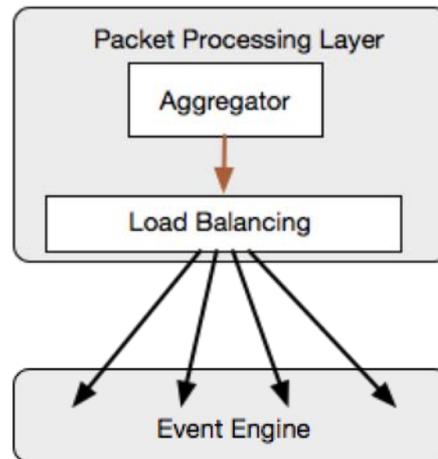
What is the Bro IDS? Packet Processing Layer

- Packet processing layer example
 - External hardware consuming and breaking out data streams to each bro node



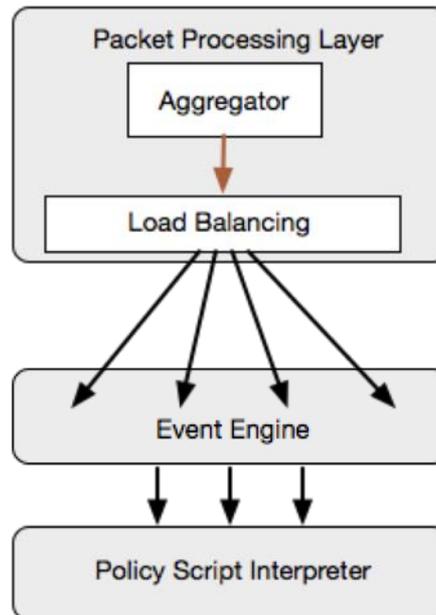
What is the Bro IDS?

- Event Engine or “Bro Core”
 - Dynamic Protocol Detection (DPD)
 - Generates “Events” to be processed



What is the Bro IDS?

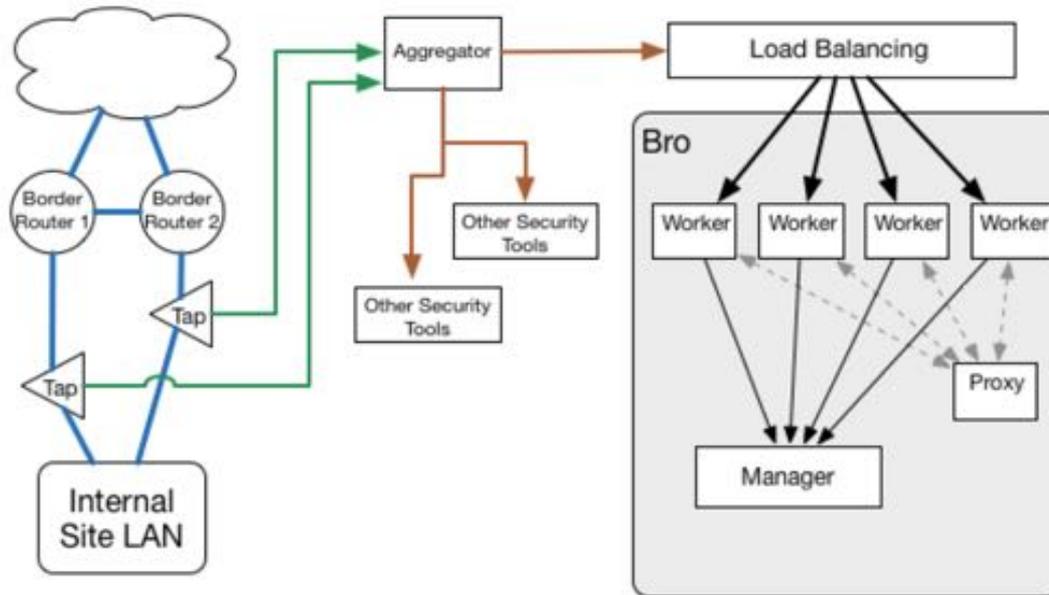
- A policy script interpreter
 - Acts on Events.
 - Bro Programming Language
 - Pre-built frameworks and protocol analyzers
 - Ships with basic policies that primarily provide logging



What does the Bro IDS do?

- Bro provides the following capabilities including (but not limited to):
 - Deep packet inspection
 - Attack and anomaly detection
 - Event correlation
 - Alert generation
 - Full IPv6 and IPv4 support
 - A powerful, flexible policy scripting language
 - Scalable, clustering architecture
- Accolades
 - Born from research and education networking
 - Used and tested in the fastest networks on the planet

Science Data and large flows



Science Data and Large Flows

- Designing Bro for 100G:
 - <http://go.lbl.gov/100g>
 - A few years old, but still relevant and quite detailed.

Integration

- Integrates into existing tools
- Utilize resources already in place
 - SIEM (Log aggregation)
 - Log hosts (Log aggregation)
 - Splunk (Log aggregation)
 - Flow data collectors (As an analog or verification tool)
 - Pagerduty (Alerting and notification)
 - custom middleware (Other proprietary services for internal process)
- Built for flexibility.
 - Scalable
 - IPv4 and IPv6 aware

Actions: Logging

```

lsnqj10@mln:/nas/bms/ops/current$ ls21 -F con.log
343298515.49838  Q33nI2K7DfDm9g 38.289.289.9 2867 204.8.8.1 8612 udp - - - - 58 T 8 0 1 44 0 8 [empty] - - onlan-eth1
343298519.38525  Q33nI2K7DfDm9g 38.289.289.7 38162 38.289.289.1 55 udp dns 0.88637 43 88 SF T 0 84 1 71 1 188 (empty) - - onlan-eth1
343298519.81988  Q33nI2K7DfDm9g 38.289.289.55 42145 239.255.255.258 2980 udp - - - - 58 T 8 0 1 122 0 8 [empty] - - onlan-eth1
343298519.88381  Q33nI2K7DfDm9g 38.289.289.5 2980 38.289.289.55 42145 udp - 0.88639 4038 0 58 T 0 1 11 4418 8 0 (empty) - - onlan-eth1
343298519.28796  Q33nI2K7DfDm9g 38.289.289.7 68989 38.289.289.1 55 udp dns 0.88667 43 88 SF T 0 84 1 71 1 188 (empty) - - onlan-eth1
343298519.28947  Q33nI2K7DfDm9g 38.289.289.7 42398 38.289.289.1 55 udp dns 0.88625 43 88 SF T 0 84 1 71 1 188 (empty) - - onlan-eth1
343298519.28958  Q33nI2K7DfDm9g 38.289.289.7 25588 38.289.289.1 55 udp dns 0.88624 43 88 SF T 0 84 1 71 1 188 (empty) - - onlan-eth1
343298519.28957  Q33nI2K7DfDm9g 38.289.289.7 44752 38.289.289.1 55 udp dns 0.88610 43 88 SF T 0 84 1 71 1 188 (empty) - - onlan-eth1
343298519.734974 Q33nI2K7DfDm9g 38.289.289.7 43655 38.289.289.1 55 udp dns 0.88702 43 88 SF T 0 84 1 71 1 188 (empty) - - onlan-eth1
343298519.30924  Q33nI2K7DfDm9g 38.289.289.9 21758 38.289.289.83 288C udp - 146.94284 9288 0 58 T 0 1 38 18018 8 0 (empty) - - onlan-eth1
343298519.40782  Q33nI2K7DfDm9g 2881.10.2880.48c-8888.1077-1ed1.288C 47858 2881.2880.84.-b 55 udp dns 0.879842 32 68 SF F 0 84 1 79 1 188 (empty) IS
IS onlan-eth1
343298519.68744  Q33nI2K7DfDm9g 38.289.289.7 26485 38.289.289.1 55 udp dns 0.88638 43 88 SF T 0 84 1 71 1 188 (empty) - - onlan-eth1
343298519.68936  Q33nI2K7DfDm9g 38.289.289.7 40162 38.289.289.1 55 udp dns 0.88670 43 88 SF T 0 84 1 71 1 188 (empty) - - onlan-eth1
343298519.798183 Q33nI2K7DfDm9g 38.289.289.9 48347 38.289.289.1 55 udp dns 0.813288 37 117 SF T 0 84 1 68 1 185 (empty) - - onlan-eth1
343298519.798183 Q33nI2K7DfDm9g 38.289.289.9 25059 38.289.289.1 55 udp dns 0.813321 37 89 SF T 0 84 1 68 1 117 (empty) - - onlan-eth1
343298521.946296 Q33nI2K7DfDm9g 38.289.289.59 41589 38.289.289.8 28228 udp - - - - 58 T 8 0 1 201 0 8 [empty] - - onlan-eth1
343298521.824279 Q33nI2K7DfDm9g 38.289.289.7 49758 38.289.289.1 55 udp dns 0.88626 43 88 SF T 0 84 1 71 1 188 (empty) - - onlan-eth1
343298521.989546 Q33nI2K7DfDm9g 38.289.289.7 42438 38.289.289.1 55 udp dns 0.88679 43 79 SF T 0 84 1 71 1 180 (empty) - - onlan-eth1
343298521.981412 Q33nI2K7DfDm9g 38.289.289.7 49675 38.289.289.1 55 udp dns 0.88625 43 79 SF T 0 84 1 71 1 180 (empty) - - onlan-eth1
343298523.83932 Q33nI2K7DfDm9g 38.289.289.8 42842 38.289.289.55 42145 udp - 0.88638 2887 0 58 T 0 1 6 2155 8 0 (empty) - - onlan-eth1
343298523.274484 Q33nI2K7DfDm9g 38.289.289.7 48414 38.289.289.1 55 udp dns 0.88654 43 88 SF T 0 84 1 71 1 188 (empty) - - onlan-eth1
343298523.274413 Q33nI2K7DfDm9g 38.289.289.7 52288 38.289.289.1 55 udp dns 0.88634 43 88 SF T 0 84 1 71 1 188 (empty) - - onlan-eth1
343298524.81418 Q33nI2K7DfDm9g 38.289.289.7 45747 38.289.289.1 55 udp dns 0.88638 43 88 SF T 0 84 1 71 1 188 (empty) - - onlan-eth1

```

Actions: Alerting

- Customizable Notification framework
- Large number of variables:
 - \$note
 - \$msg
 - \$sub
 - \$conn
 - \$id
 - \$src
 - \$n
 - \$identifier
 - \$suppress_for



Out of the box....

- Connection Log
 - Similar to netflow information
- Protocol specific logs:
 - HTTP, FTP, SMTP, IRC, SSH, SSL, DNS, ...
- Observational logs:
 - known_certs, known_services, known_devices, software, files
- Detection:
 - Intel, notice, notice_alarm, signatures, traceroute
- Diagnostics
 - capture_loss, packet_filter, communication, reporter

Out of the box: Notices

CaptureLoss::Too_Much_Loss

Conn::Retransmission_Inconsistency

Conn::Ack_Above_Hole

Conn::Content_Gap

DNS::External_Name

FTP::Bruteforcing

FTP::Site_Exec_Success

HTTP::SQL_Injection_Attacker

HTTP::SQL_Injection_Victim

Heartbleed::SSL_Heartbeat_Attack

Heartbleed::SSL_Heartbeat_Attack_Success

Heartbleed::SSL_Heartbeat_Odd_Length

Heartbleed::SSL_Heartbeat_Many_Requests

Intel::Notice

PacketFilter::Compile_Failure

PacketFilter::Install_Failure

PacketFilter::Too_Long_To_Compile_Filter

PacketFilter::Dropped_Packets

PacketFilter::Cannot_BPF_Shunt_Conn

ProtocolDetector::Protocol_Found

ProtocolDetector::Server_Found

SMTP::Blocklist_Error_Message

SMTP::Blocklist_Blocked_Host

SMTP::Suspicious_Origination

SSH::Password_Guessing

SSH::Login_By_Password_Guesser

SSH::Watched_Country_Login

SSH::Interesting_Hostname_Login

SSL::Certificate_Expired

SSL::Certificate_Expires_Soon

SSL::Certificate_Not_Valid_Yet

SSL::Invalid_Server_Cert

SSL::Invalid_Ocsp_Response

SSL::Weak_Key

SSL::Old_Version

SSL::Weak_Cipher

Scan::Address_Scan

Scan::Port_Scan

Signatures::Sensitive_Signature

Signatures::Multiple_Signatures

Signatures::Multiple_Sig_Responders

Signatures::Count_Signature

Signatures::Signature_Summary

Software::Software_Version_Change

Software::Vulnerable_Version

Traceroute::Detected

Weird::Activity

Care and feeding

- Consume community intelligence feeds
 - Alert based on a combination of criteria from different feeds
- Tuning, tuning, tuning.
 - Not “set and forget”**
 - There exists a large and active community of Really Smart People writing Bro policies

**** Out of the box, untuned Bro IDS will still provide huge amounts of useful information.**

Actions

- Execute external scripts for operational response
 - Black hole routing
 - <https://github.com/buraglio/singularity>
 - <https://github.com/JustinAzoff/bhr-site>
 - Apply ACLs
 - Quarantine hosts
 - ...basically anything that you can write a script to do

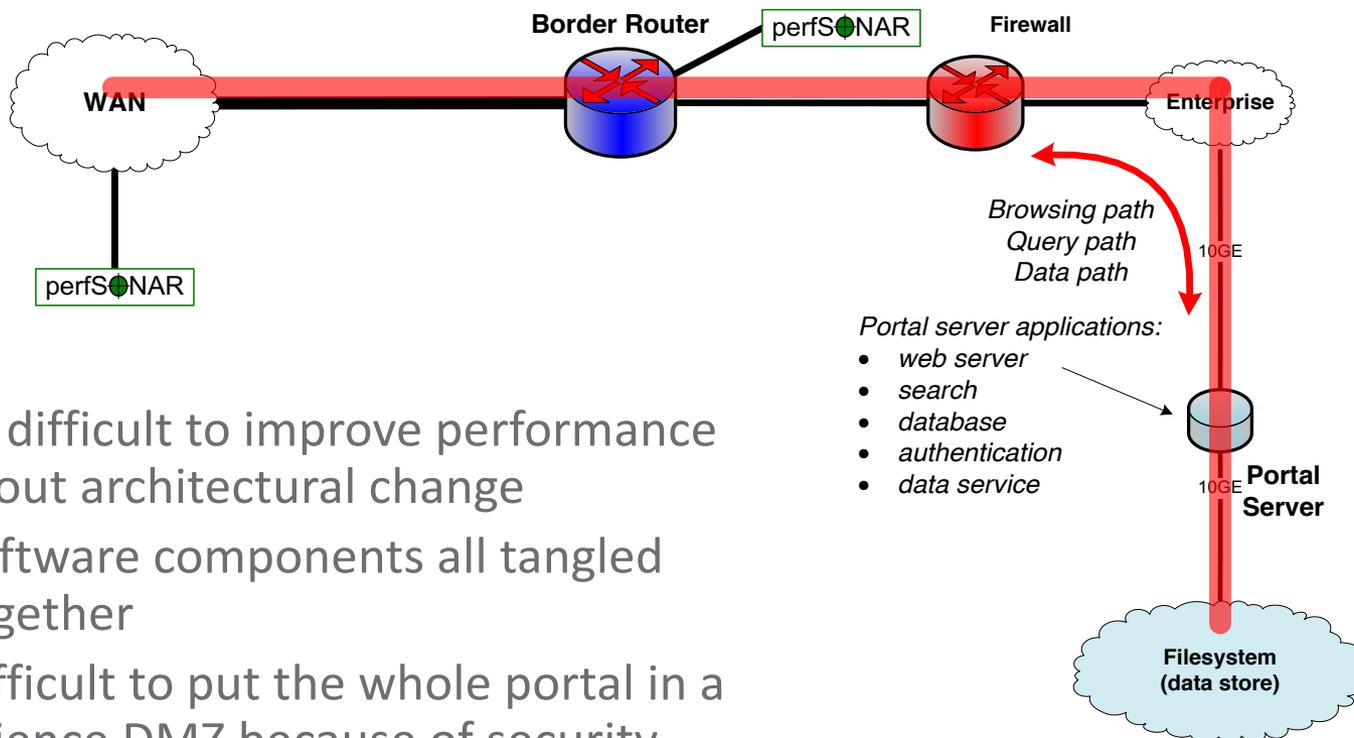
Signature-based IDS

- Many networks today run both Bro and a signature based IDS.
- Suricata appears to be favored, but Snort used to be.
- That pendulum may swing again, but the Bro pendulum hasn't.
- Signature-based systems often used with intel/threat feeds (e.g. Emerging Threats).
- Bro can also make use of feeds.

Science Data Portals

- Large repositories of scientific data
 - Climate data
 - Sky surveys (astronomy, cosmology)
 - Many others
 - Data search, browsing, access
- Many scientific data portals were designed 15+ years ago
 - Single-web-server design
 - Data browse/search, data access, user awareness all in a single system
 - All the data goes through the portal server
 - In many cases by design
 - E.g. embargo before publication (enforce access control)
 - Better than old command-line FTP, but outdated by today's standards

Legacy Portal Design

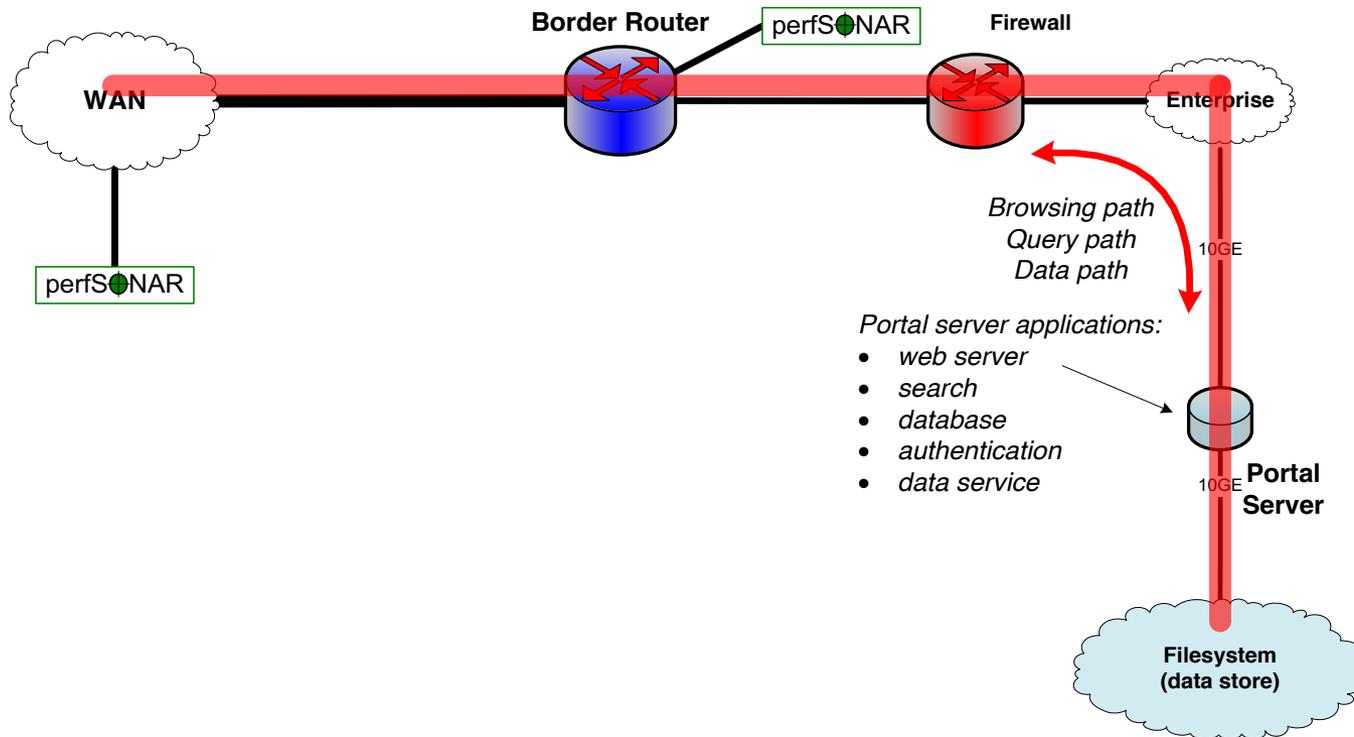


- Very difficult to improve performance without architectural change
 - Software components all tangled together
 - Difficult to put the whole portal in a Science DMZ because of security
 - Even if you could put it in a DMZ, many components aren't scalable
- What does architectural change mean?

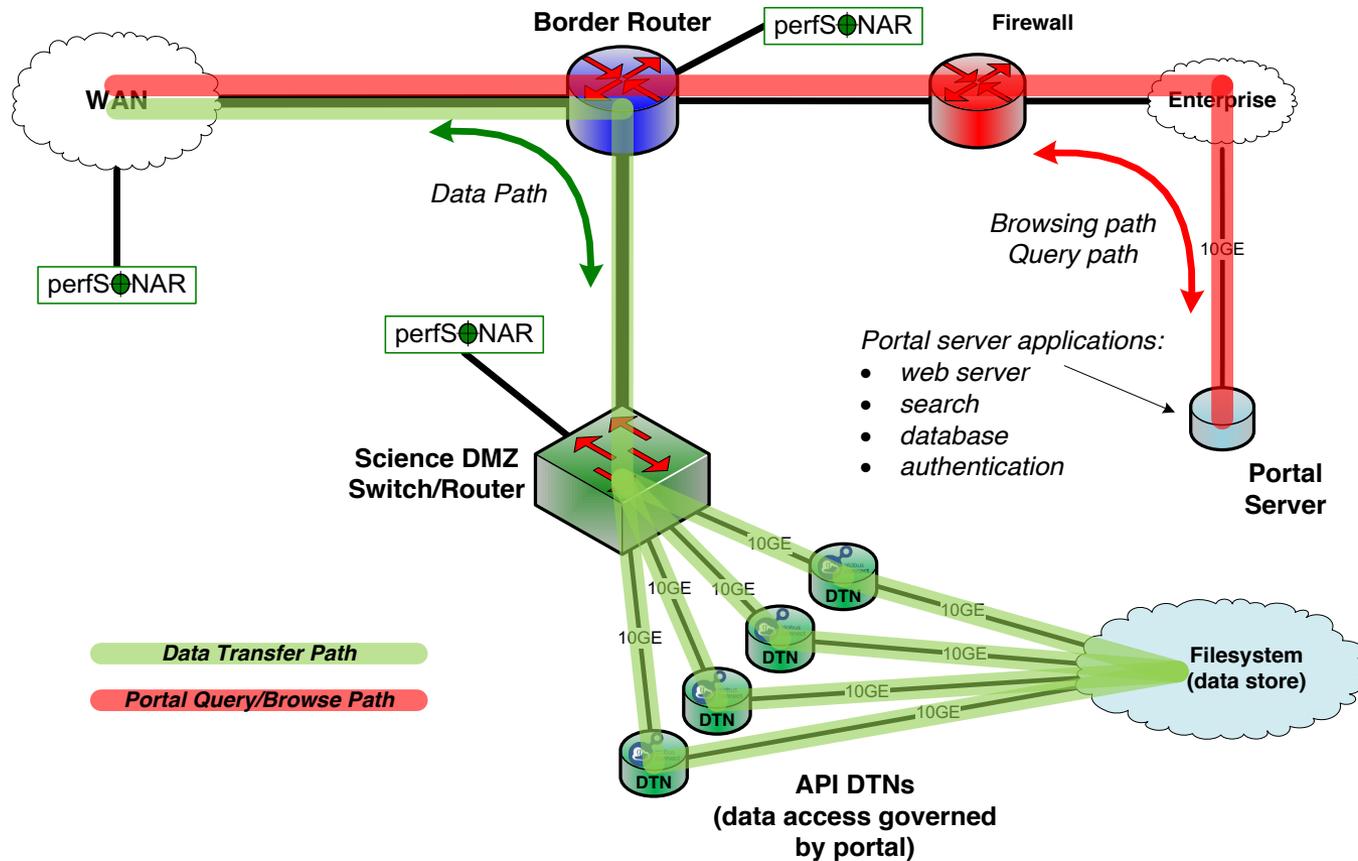
Architectural Examination of Data Portals

- Not necessarily advocating CDNs for science data (not really a good fit)
- Common data portal functions (most portals have these)
 - Search/query/discovery
 - Data download method for data access
 - GUI for browsing by humans
 - API for machine access – ideally incorporates search/query + download
- Performance pain is primarily in the data handling piece
 - Rapid increase in data scale eclipsed legacy software stack capabilities
 - Portal servers often stuck in enterprise network
- Can we “disassemble” the portal and put the pieces back together better?
 - Use Science DMZ as a platform for the data piece
 - Avoid placing complex software in the Science DMZ

Legacy Portal Design



Next-Generation Portal Leverages Science DMZ



<https://peerj.com/articles/cs-144/>

Put The Data On Dedicated Infrastructure

- We have separated the data handling from the portal logic
- Portal is still its normal self, but enhanced
 - Portal GUI, database, search, etc. all function as they did before
 - Query returns pointers to data objects in the Science DMZ
 - Portal is now freed from ties to the data servers (run it on Amazon if you want!)
- Data handling is separate, and scalable
 - High-performance DTNs in the Science DMZ
 - Scale as much as you need to without modifying the portal software
- Outsource data handling to computing centers
 - Computing centers are set up for large-scale data
 - Let them handle the large-scale data, and let the portal do the orchestration of data placement
- <https://peerj.com/articles/cs-144/> - Modern Research Data Portal paper

NCAR RDA Data Portal

- Let's say I have a nice compute allocation at NERSC – climate science
- Let's say I need some data from NCAR for my project
- <https://rda.ucar.edu/>
- Data sets (there are many more, but these are two):
- <https://rda.ucar.edu/datasets/ds199.1/>
- <https://rda.ucar.edu/datasets/ds313.0/>
- Download to NERSC (could also do ALCF or NCSA or OLCF)

NCAR's Research Data Archive X
 Secure | https://nda.ucar.edu

UCAR - NCAR Closure/Emergencies Locations/Directions Find People

Hello dart@es.net dashboard sign out

NCAR UCAR | **Research Data Archive**
 Computational & Information Systems Lab *weather • data • climate*

Go to Dataset:

Home Find Data Ancillary Services About/Contact Data Citation Web Services For Staff

First-time visitor to our site?
 Please take a video tour of our home page

Dataset Search:
 Search Advanced Options

Look For Data By:

All Datasets	Variable/Parameter	Type of Data
Time Resolution	Platform	Spatial Resolution
Topic/Subtopic	Project/Experiment	Supports Project
Data Format	Instrument	Location
	Recently Added/Updated	

Recently Added Datasets: (within the last 6 months)

- ERAS Reanalysis Monthly Means
- Daily Gridded North American Snowfall
- ERAS Reanalysis
- NCAR/MOPITT Reanalysis
- GridRad - Three-Dimensional Gridded NEXRAD WSR-88D Radar Data
- CMIP 5 dataset and code for R parallelization
- Dai and Trenberth Global River Flow and Continental Discharge Dataset
- Dai Global Palmer Drought Severity Index (PDSI)

Get Help:

- Frequently Asked Questions
- Reset your password
- A-Z Site Index
- RDA Users Email List
- RDA Blog
- RDA video tutorials
- Email Us

From Our Blog:

- Accessing RDA OPeNDAP endpoints with authentication
- All RDA data transfer and processing services restored to production
- RDA Service Outage July 14-18, 2017
- RDA web services down for maintenance at 1PM MDT on May 3, 2017

[More blog posts ...](#)

GLADE Users:
 Much of the RDA is directly accessible from CSI's Globally Accessible Data

NCAR's Research Data Archive X

Secure | <https://rda.ucar.edu/#?id?nb=y&b=a&v=Full+List>

UCAR | NCAR

Closures/Emergencies | Locations/Directions | Find People

Hello [dart@es.net](#) | [dashboard](#) | [sign out](#)

NCAR |  **Research Data Archive**
Computational & Information Systems Lab

weather • data • climate

Go to Dataset:

Home | Find Data | Ancillary Services | About/Contact | Data Citation | Web Services | For Staff

Look For Data:

- Create a New List
- OR --
- Continue Narrowing By:
- Variable / Parameter
- Type of Data
- Time Resolution
- Platform
- Spatial Resolution
- Topic / Subtopic
- Project / Experiment
- Supports Project
- Data Format
- Instrument
- Location
- Progress
- Free Text

Browse the RDA

Showing datasets with these attributes: [All RDA Datasets](#) | [Full List](#) (180)

Select two datasets and [Compare](#) them. [Reset](#) checkboxes

[1. Daily Northern Hemisphere Sea Level Pressure Grids, continuing from 1899 \(ds010.0\)](#)

grids contained in this dataset make up the longest continuous set of daily gridded pressure data in the D55 archive. These grids have been ... [\(i\)](#)

[where Sea-Level Pressure Grids, continuing from 1899 \(ds010.1\)](#)

continuous time series of monthly gridded Northern Hemisphere sea-level pressure degree latitude/longitude grids, computed from the daily grids in ... [\(i\)](#)

[where Daily Sea-Level Pressure Grids for 1880 to 1979 \(ds012.0\)](#)

Northern Hemisphere sea-level pressure data on a 10-degree by 5-degree (36x16) grid period 1880 to 1979.

[where Daily \(and Monthly\) Sea-Level Pressure and 500 mb Height Grids for 1946Jan to 1993Dec \(ds018.0\)](#)

The gridded daily sea-level pressure analyses in this dataset were produced by the operational models of the U.S. Navy Fleet Numerical Oceanography Center (FNO). The data are arranged in a ... [\(i\)](#)

Type of Data
• Climate_proxy (2)
• Grid (389)
• Log_entry (1)
• Platform_observation (258)
• Radar (2)
• Satellite (38)



GEOSS Global Atmosphere Forcing Data

ds313.0 ☆

For assistance, contact Chi-Fan Shih (303-497-1833).

Description

Data Access

Help with this page: [RDA dataset description page video tour](#)

Abstract: GEOSS Atmospheric Forcing data, regridded and prepared as meteorological variables to run CESM and WRF simulations.

Temporal Range: 2004-01-02 00:00 +0000 to 2017-10-19 21:00 +0000 (Entire dataset)

↑ [Period details by dataset product](#)

Updates: Irregularly

Variables: [Surface Pressure](#) [Upper Level Winds](#)

↑ [Variables by dataset product](#)

Vertical Levels: See the detailed [metadata](#) for level information

Data Types: Grid

Spatial Coverage: Longitude Range: Westernmost=180W Easternmost=180E

Latitude Range: Southernmost=90S Northernmost=90N

↑ [Detailed coverage information](#)

Data Contributors: [UCAR/NCAR/ACD](#) | [UCAR/NCAR/CGD](#)

How to Cite This Dataset:

[RIS](#)
[BibTeX](#)

Tilmes, S., 2016. *GEOSS Global Atmosphere Forcing Data*. Research Data Archive at the National Center for Atmospheric Research, Computational and Information Systems Laboratory. <http://rda.ucar.edu/datasets/ds313.0/>. Accessed[†] dd mmm YYYY.

[†]Please fill in the "Accessed" date with the day, month, and year (e.g., 5 Aug 2011) you last accessed the data from the RDA.

Bibliographic citation shown in [Federation of Earth Science Information Partners \(ESIP\)](#) style

[Get a customized data citation](#)

Total Volume: 449.28 GB

Data Formats: *netCDF*

More Details: [View more details](#) for this dataset, including dataset citation, data contributors, and other detailed metadata

Data Access: Click [Data Access](#) here or in the navigation bar near the top of the page

Metadata Record: Display in [choose from the list](#) format

UCAR NCAR | Hello dart@es.net | dashboard | sign out

NCAR UCAR | **Research Data Archive** | Computational & Information Systems Lab | *weather • data • climate*

Go to Dataset:

Home | Find Data | Ancillary Services | About/Contact | Data Citation | Web Services | For Staff

GEOSS Global Atmosphere Forcing Data

ds313.0 ☆

For assistance, contact Chi-Fan Shih (303-497-1833).

Description | **Data Access**

Mouse over the table headings for detailed descriptions

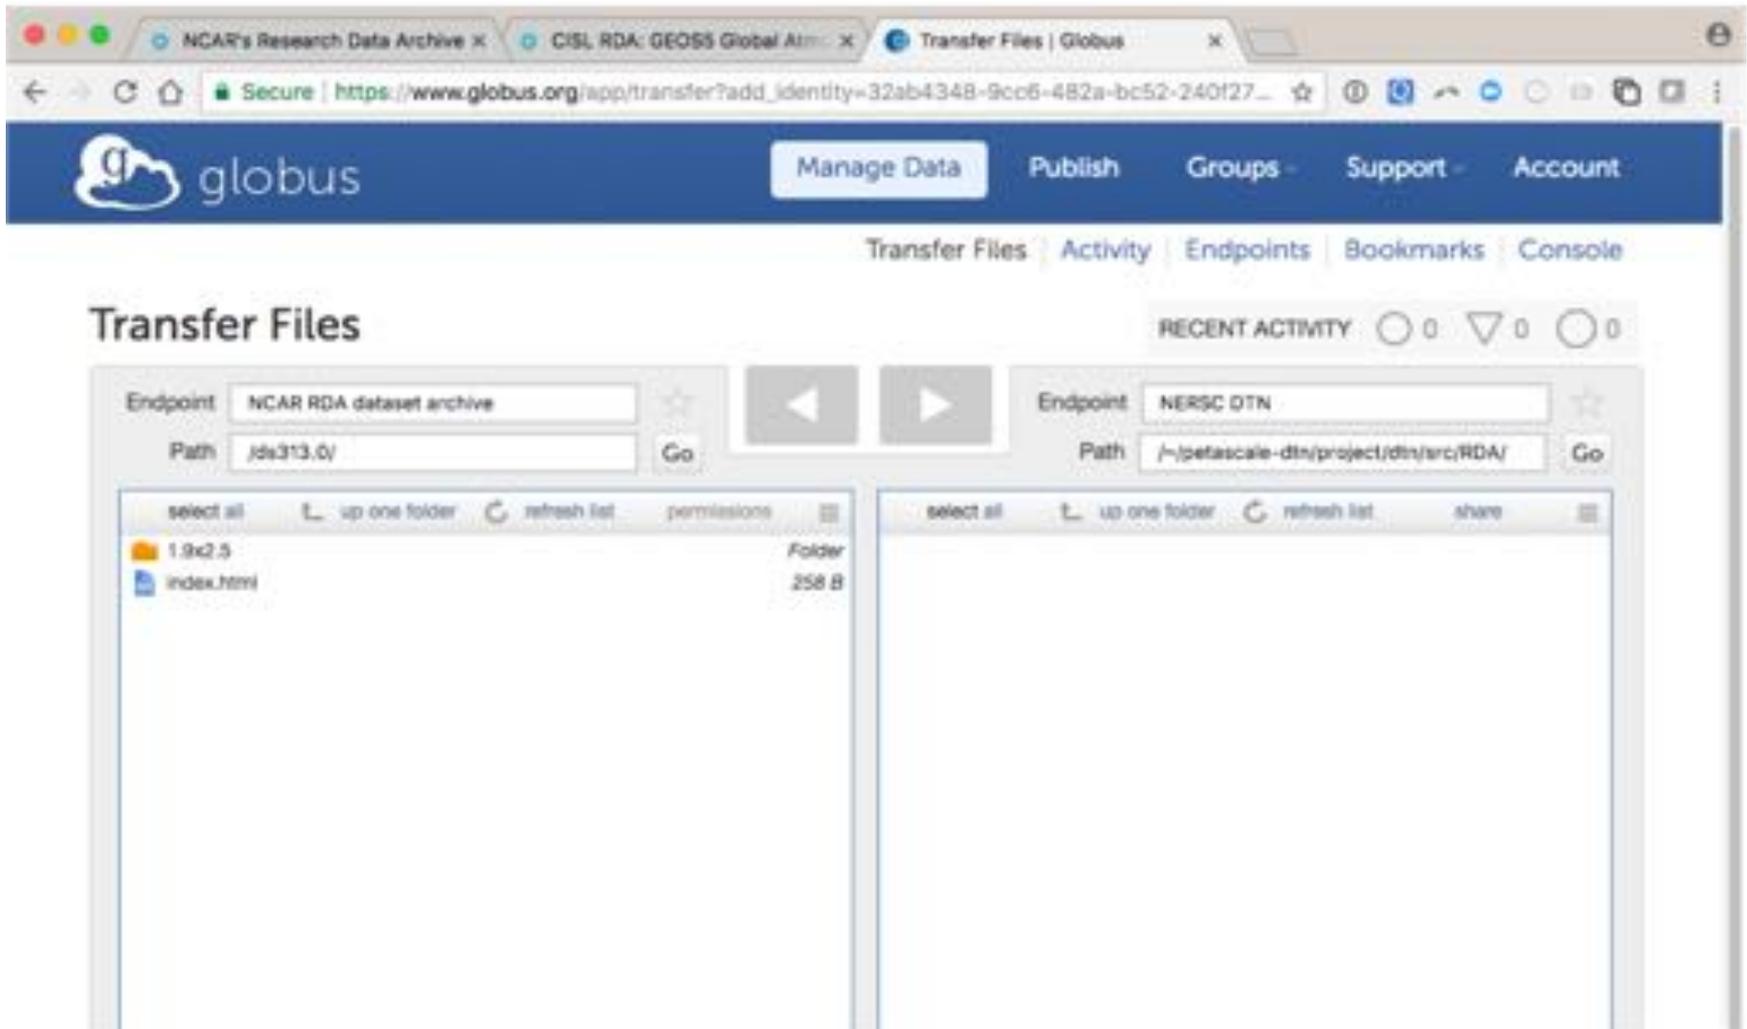
Data File Downloads		NCAR-Only Access	
Web Server Holdings	Globus Transfer Service (GridFTP)	Central File System (GLADE) Holdings	Tape Archive (HPSS) Holdings
Web File Listing	Globus Transfer	GLADE File Listing	HPSS File Listing

The Research Data Archive is managed by the Data Support Section of the Computational and Information Systems Laboratory at the National Center for Atmospheric Research in Boulder, Colorado. NCAR is sponsored by the National Science Foundation.

Follow us: Atom Facebook Twitter

© 2017, UCAR | [Privacy Policy](#) | [Terms of Use](#) | [Contact Us](#)

Portal creates a Globus transfer job for us



Submit the transfer job, go about our business

The screenshot displays the Globus Transfer Files web interface. At the top, the Globus logo is on the left, and navigation links for 'Manage Data', 'Publish', 'Groups', 'Support', and 'Account' are on the right. Below this, a secondary navigation bar includes 'Transfer Files', 'Activity', 'Endpoints', 'Bookmarks', and 'Console'. The main heading is 'Transfer Files', with a 'RECENT ACTIVITY' section showing 1 success, 0 failures, and 0 pending items. A green notification banner at the top of the main content area states: 'Transfer request submitted successfully. Task id: d2776d02-bb8f-11e7-9428-22000a8cbd7d'. Below the notification, two endpoint panels are visible. The left panel shows the source endpoint 'NCAR RDA dataset archive' with path '/ds3/3.0/'. The right panel shows the destination endpoint 'NERSC DTN' with path '/-/petascale-dtn/project/dtn/src/RDA/'. Between the panels are navigation arrows. Below each endpoint panel is a file list. The left list shows a folder '1.8x2.5' and a file 'index.html'. The right list is currently empty.

Data Transfer from RDA Portal – Results

Activity

Task List



NCAR RDA dataset archive to NERSC DTN

transfer completed 5 hours ago



Overview



Event Log

Task ID 4f923e48-bb48-11e7-9428-22000a8cbd7d

Owner Eli Dart (dart@globusid.org)

Source NCAR RDA dataset archive ⓘ
owner: rda@globusid.org

Destination NERSC DTN ⓘ
owner: nersc@globusid.org

Condition SUCCEEDED

Requested 2017-10-27 11:54 am

Completed 2017-10-27 11:58 am

Transfer Settings

- verify file integrity after transfer
- transfer is not encrypted
- overwriting all files on destination

Files 5041

Directories 15

Bytes Transferred 449.27 GB

Effective Speed 1.84 GB/s

Pending 0

Succeeded 5057

Cancelled 0

Expired 0

Failed 0

Retrying 0

Skipped 0

[view debug data](#)

Conclusion, wrap up

- The Science DMZ is an integral part of a good risk-based security posture for the entire institution.
- Work with security teams, not against them—a well-designed Science DMZ is good for them.
- Think about *practical* risks and design controls around them.
- Tailor the security solutions to the *thing* you are trying to secure!
- Don't forget about function!
- Intrusion detection and packet protection can scale to a small number of 100GE links.