



FUNCTIONAL SECURITY AT TACC

TACC AT A GLANCE



Personnel

160 Staff (~70 PhD)

Facilities

12 MW Data center capacity
Two office buildings, Three
Datacenters, two visualization
facilities, and a chilling plant.

Systems and Services


A Billion compute hours per year
5 Billion files, 50 Petabytes of Data,
Hundreds of Public Datasets

Capacity & Services

HPC, HTC, Visualization, Large scale
data storage, Cloud computing
Consulting, Curation and analysis,
Code optimization, Portals and
Gateways, Web service APIs, Training
and Outreach



EXTREME SCALE SUPERCOMPUTING



Stampede

- #10 HPC system in the world for computation 500k CPU core 9.7 PF

Lonestar 5

- Texas-focused Cray XC40 30,000 Intel Haswell cores 1.25 PF

Wrangler

- 0.6 PB usable DSSD flash storage w 1 TB/s read rate + 10 PB Lustre

Maverick

- 132 Fat nodes w dual 10 core Ivy Bridge + NVIDIA Kepler K40 GPGPU

Chameleon & Jetstream Cloud

- 1400 nodes OpenStack

Disk and Tape Storage

- 100+ PB storage in HIPAA-aligned data center



Hikari

- 380V DC Green computing system partnership with NEDO and NTT. 10k Haswell cores. HVDC and Solar (partial)
- Support for container ecosystem

ASSUMPTIONS

- ▶ ~5% of user accounts are usually compromised (from things they have done elsewhere)
- ▶ Users make poor choices (almost always)
- ▶ Sysadmins might also make poor choices (less often, but it happens).
- ▶ State sponsored attacks will occasionally succeed (they have lots of resources)
- ▶ You can't stop them all
- ▶ No one is perfect (including you)

CHANGING PATTERNS OF USE

- ▶ SSH Users: Decreasing Yearly (but still a whole lot, and still indispensable)
- ▶ Portal Users: Increasing Yearly
- ▶ API's: Increasing Yearly
- ▶ VM's and Containers: The future
 - ▶ Open Stack: Jetstream (IU and TACC) and Chameleon (UC/ANL and TACC).
 - ▶ VM's
 - ▶ Bare Metal
 - ▶ SDN
 - ▶ Docker: Developers Best Friend
 - ▶ No real security
 - ▶ Black Boxes

ONE SIZE DOES NOT FIT ALL

- ▶ “Classic” users love SSH and haven't changed or will change
- ▶ Portal Users don't ever use SSH but only use the web gui's
- ▶ Cloud Users spin up VM's and may need root
- ▶ API users hit many resources via API and enable workflows
- ▶ Container Users come in all of the above
- ▶ All those users need different types of security and network configurations

SECURITY BASICS

- ▶ Patch, patch, patch, patch, and patch again
- ▶ Log, Log, Log, Log – everything you can, as much as you can, and keep it forever.
- ▶ Lockout policy (bad stuff happens from abandoned accounts)
- ▶ Employee Checkout – have procedures in place when someone leaves
- ▶ IDS/NSM (Bro) – Have one!
- ▶ Scan systems – from inside and out.

BASICS(2)

- ▶ MFA (RSA) for Admins
- ▶ Sudo/LUP -- Don't give out more privilege than you need to.
- ▶ Keep it simple
 - ▶ Overcomplicating leads to users/staff not doing the right thing
- ▶ Make the easy choices
- ▶ Read Only Friday (everyone loves this)
- ▶ Staff Development

CULTURE

- ▶ Modify culture (giving root is not a right)
- ▶ Stick of Compliance (use it to modify old/less insecure practices)
- ▶ Have everyone participate (including your gray beards)
 - ▶ Sysadmins will be on board if they know the goal
- ▶ Management buy in (you need it to succeed)
- ▶ Talk about it (and more after that)
- ▶ Teach new users/staff good habits
- ▶ Staff trainings

RE-THINK OLD IDEAS

- ▶ Circle the wagons is outdated
- ▶ In CI, everything is a DMZ
- ▶ Deep forensics have limited value
 - ▶ Do enough to know how they got in
- ▶ Re-think metrics
- ▶ Re-think success

METRICS @ TACC

- ▶ Metrics we like
 - ▶ Time to detection (Should be sub 10 min)
 - ▶ Time to resolution (Should be sub 30 min)
 - ▶ Number of failed login attempts (If above the baseline something is up)
 - ▶ Data movement (If above the baseline something is up)
 - ▶ Number of attacks (If above the baseline something is up)
 - ▶ Number of actionable events (Did you actually have to do anything)
 - ▶ Intrusions

LET YOUR CAMPUS DO FOR YOU THE THINGS THEY ARE GOOD AT!

- ▶ Email (no you don't need to run your own email)
- ▶ Box
- ▶ Stache (Secure information sharing)
- ▶ Building Access and Control Systems (BACS)
- ▶ Video Surveillance
- ▶ Physical Security (Campus PD)

INVENTORY TOOLS @ TACC

- ▶ You can't protect what you don't know about
- ▶ DCIM (Data Center Information Management)
 - ▶ Data Center Map
- ▶ Dopplr (IPAM)
 - ▶ phpIPAM
- ▶ Solar Winds

CLOUDS AND CONTAINERS



- ▶ The Future!
- ▶ You are a Service Provider
- ▶ More threats than we imagined
- ▶ Different landscape
 - ▶ No direct management of VM's
- ▶ Black Boxes
- ▶ Developed tools to find and terminate bad VM's and containers
 - ▶ Done by 24hrs operations team



PHYSICAL SECURITY @ TACC

- ▶ Visitor Policy Posted to Staff Wiki and Communicated to Staff
- ▶ All Visitors Must Check In
- ▶ Cameras on all doors
 - ▶ Monitored by operations team
 - ▶ Uses Campus Vetted System (PD Blessed)
- ▶ Maximum Number of Visitors per tour host
- ▶ ~3000 Visitors per year

CENTRAL CONFIGURATION MANAGEMENT

- ▶ Notifications when changes happen
- ▶ Central Management
- ▶ Auditable
- ▶ Watch Closer
- ▶ Master Nodes
 - ▶ RSA'd
 - ▶ LSOF
 - ▶ Central logging of cluster
 - ▶ Master Node then sends all cluster logs to Splunk
- ▶ Everything else runs Puppet or Ansible
 - ▶ RSA'd
- ▶ Solarwinds



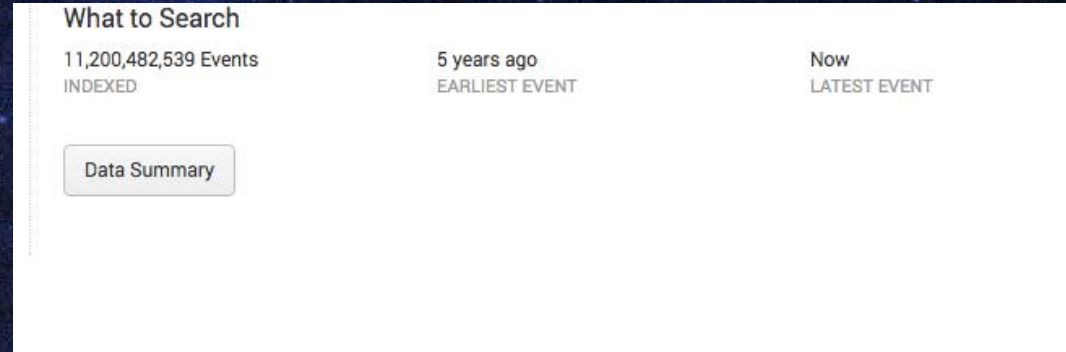
MONITORING/SCANNING @ TACC

- ▶ Monitoring and scanning will find issues
- ▶ Nagios (systems)
- ▶ Solarwinds (network)
- ▶ Splunk (everything)
- ▶ Rapid7(external)

The Nagios logo consists of the word "Nagios" in a bold, black, sans-serif font. The letter "N" is underlined. A registered trademark symbol (®) is located at the top right of the word. The logo is set against a white rectangular background.The Rapid7 logo features the word "RAPID" in a bold, black, sans-serif font, followed by a large, stylized number "7" in orange. The logo is set against a white rectangular background.

SPLUNK

- ▶ Combine all sources of information
- ▶ Smarter searching = faster results
- ▶ Granular permissions
- ▶ Expensive yes but worth it to us
- ▶ You can also use Elastic Stack



The screenshot shows a Splunk search interface. At the top, it displays 'What to Search' with a count of '11,200,482,539 Events INDEXED'. Below this is a 'Data Summary' button. To the right, there are two filters: '5 years ago EARLIEST EVENT' and 'Now LATEST EVENT'.

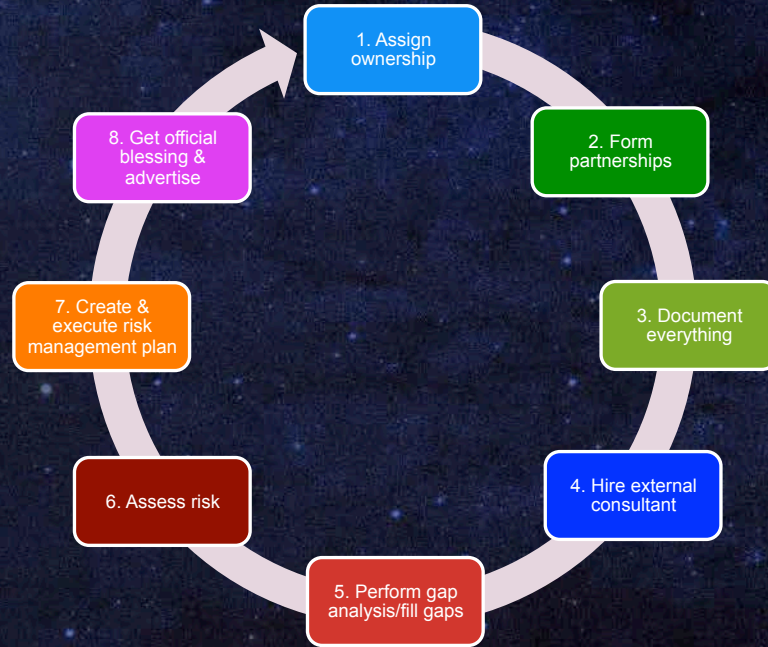
What to Search	5 years ago	Now
11,200,482,539 Events INDEXED	EARLIEST EVENT	LATEST EVENT

Data Summary

COMPLIANCE

- ▶ TACC Currently accepts HIPAA, FERPA, FISMA (Moderate), ITAR, EAR and others
- ▶ Pick a controls framework
 - ▶ TACC uses NIST
- ▶ Be willing to modify long standing policy's to meet compliance
- ▶ Continuous monitoring is key
- ▶ Be ready to write a lot
- ▶ Devote enough resources (FTE)
- ▶ Use a project management tool for tracking (redmine, JIRA, etc.)

HIPAA IMPLEMENTATION STEPS



DO WHAT YOU SAID YOU WOULD DO

- ▶ The key to all compliance is to **actually** do what you said (or documented) you would
- ▶ Verify that you actually did what you said you would
 - ▶ If you said you would do something quarterly make sure you do
- ▶ Have a third party **verify** that you did what you said you would
- ▶ Audits are not fun
 - ▶ But: they make sure you are doing what you think you are.

SOME TIMES YOU'VE GOT TO BUILD YOUR OWN

- ▶ MFA
- ▶ SSHD (need iSSHD and HPN)
- ▶ LOSF
- ▶ Puppet
- ▶ Dopplr (phpIPAM)

MFA @ TACC

- ▶ Had partnership with Toopher for all users
 - ▶ Toopher acquired
 - ▶ Backed out partnership
- ▶ Evaluated others
 - ▶ Duo, RSA, Gemalto, Yubikey
 - ▶ All were cost prohibitive
- ▶ LINOTP
- ▶ Wrote own apps (apple & android)
- ▶ SMS support
- ▶ Sourced our own hard tokens
 - ▶ Users are charged a modest fee for hard tokens
 - ▶ Soft and SMS tokens are free
- ▶ Still need RSA for root level privileges



A WORD ABOUT FIREWALLS

- ▶ Yes you need them
- ▶ Allows black hole routing
- ▶ Central administration
- ▶ Admins can self service their firewall needs @ TACC
 - ▶ Address books
- ▶ No local firewalls unless authorized by security teams @ TACC
 - ▶ In Puppet/Ainsible/etc
- ▶ Worth the money

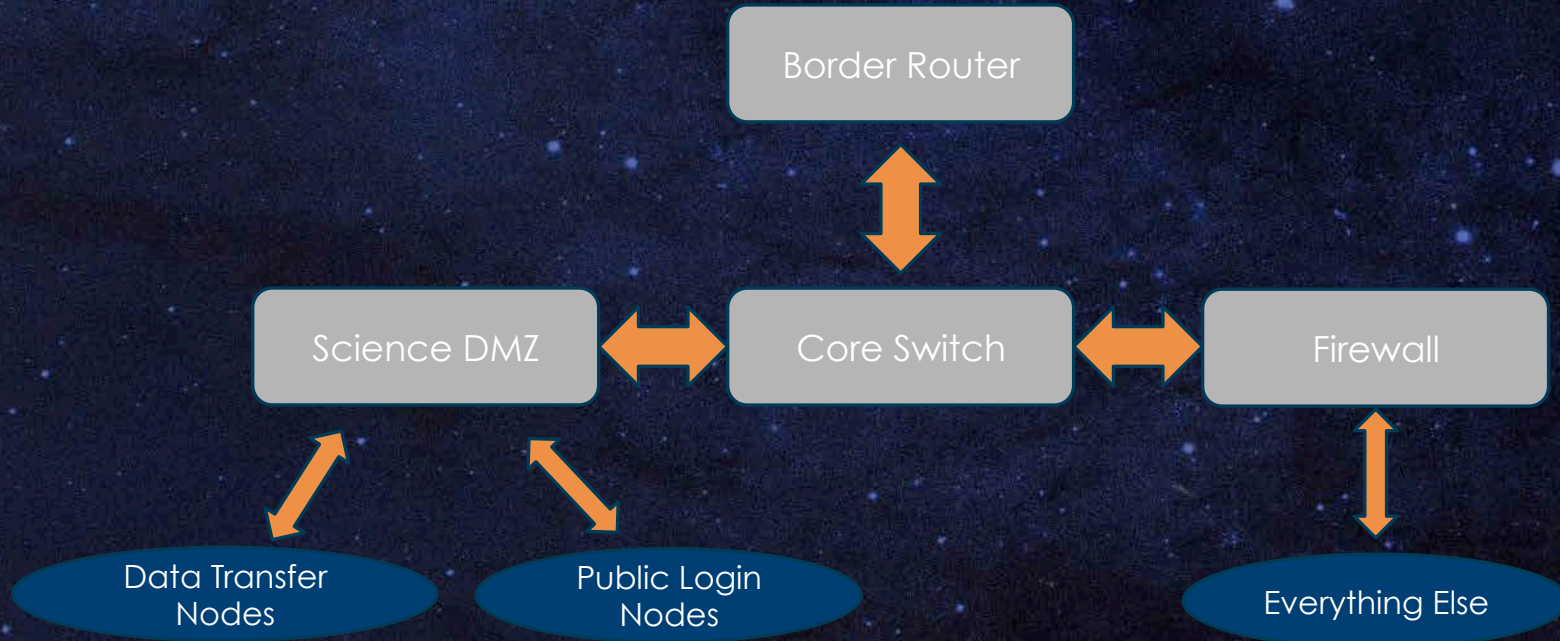
FUTURE FIREWALL @ TACC

- ▶ The network is the firewall and the firewall is the network
 - ▶ Moving from Monitoring north-south traffic to including east-west traffic
- ▶ Future firewalls will be distributed and virtualized
 - ▶ SDN based
 - ▶ Virtualized/Container based
 - ▶ Dis-aggregated hardware
 - ▶ Run on merchant silicon

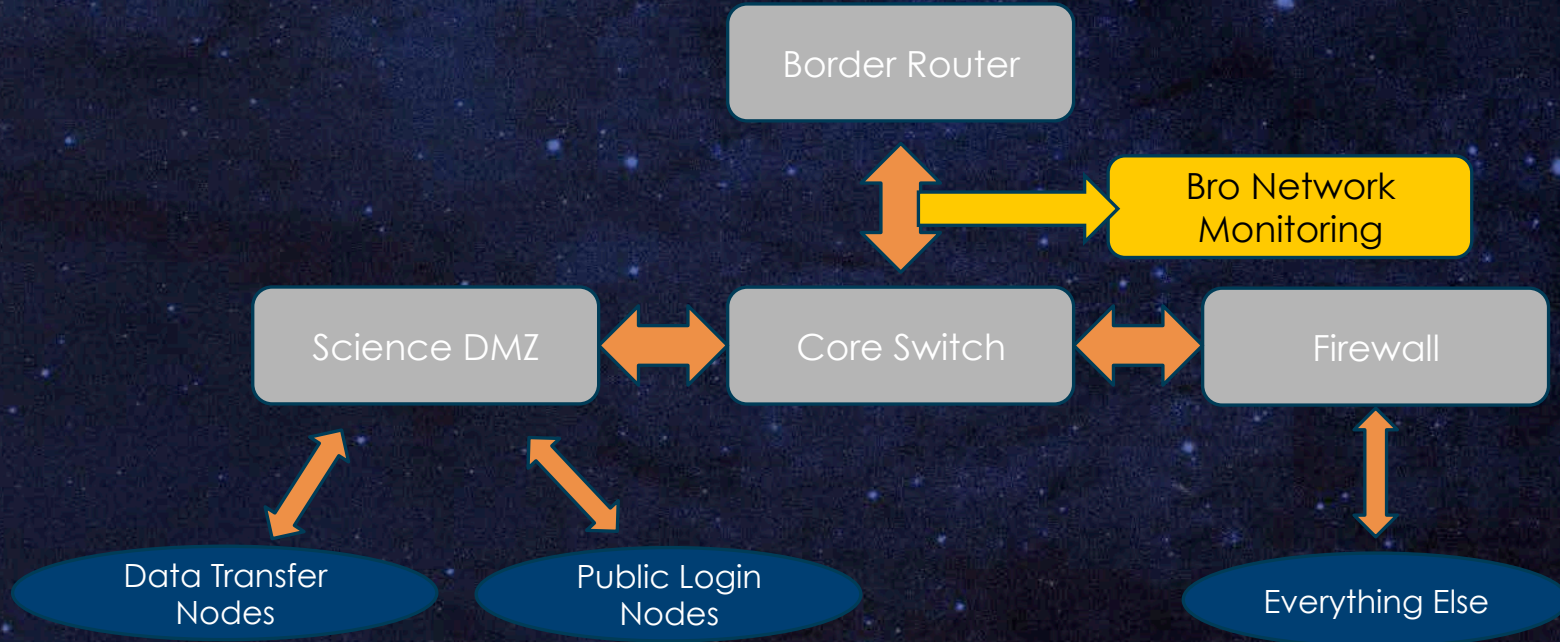
SUMMARY OF TOOLS @ TACC

- ▶ Bro
- ▶ Splunk
- ▶ SolarWinds
- ▶ phpIPAM
- ▶ OpenDCIM (to be replaced)
- ▶ Puppet
- ▶ Ainsible
- ▶ LSOF
- ▶ LINOTP
- ▶ HEAT LANrev
- ▶ FireAMP
- ▶ Redmine
- ▶ JIRA
- ▶ RT
- ▶ RSA
- ▶ Slack
- ▶ Stache
- ▶ Rapid7
- ▶ NFSEN NFDUMP
- ▶ Envoy


NETWORK MAP




NETWORK MAP (2)




IDS



IDS



IDS



IDS

