LA-UR-20-26678
Approved for public release; distribution is unlimited.

Title: HPC Infrastructure Security - Ensuring Storage and Network Safety, Integrity, Efficiency, and Performance

Author(s): Martinez, Jesse Edward
Connor, Carolyn Marie
Hollander, Brett Jason
Paschke, Kierstyn Michelle

Intended for: Virtual Conference - Tapia 2020

Issued: 2020-08-27


Los Alamos
NATIONAL LABORATORY
EST. 1943

HPC Infrastructure Security

Ensuring Storage and Network Safety, Integrity, Efficiency, and Performance

Tapia Conference 2020

Presenters:

Carolyn Connor
Brett Hollander
Jesse Martinez
Kierstyn Paschke

September 17th, 2020

Agenda

Trinity Supercomputer at Los Alamos National Laboratory

- LANL HPC Introduction
- Information Security Matters
- Network (The Nervous System - Tying it All Together)
- Data Matters: Data at Rest and Data in Motion
- Exploring Alternative and Complementary Approaches
- Hands on!

HIGH PERFORMANCE COMPUTING

What is a National Laboratory?

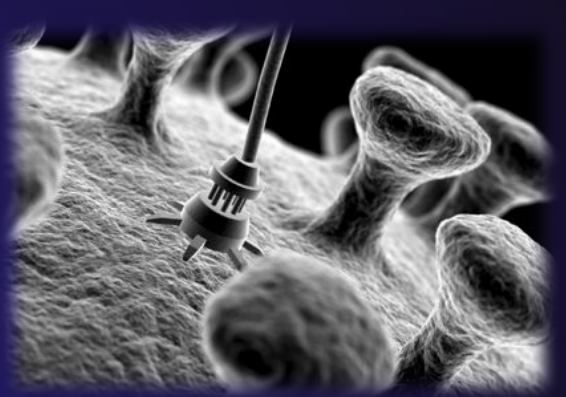

Science that matters in the national interest

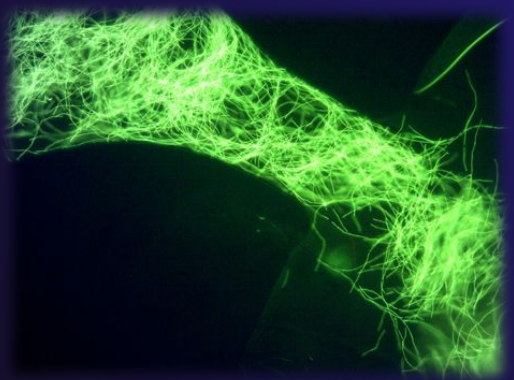

Where is Los Alamos National Laboratory ?

What is there to do around Los Alamos?

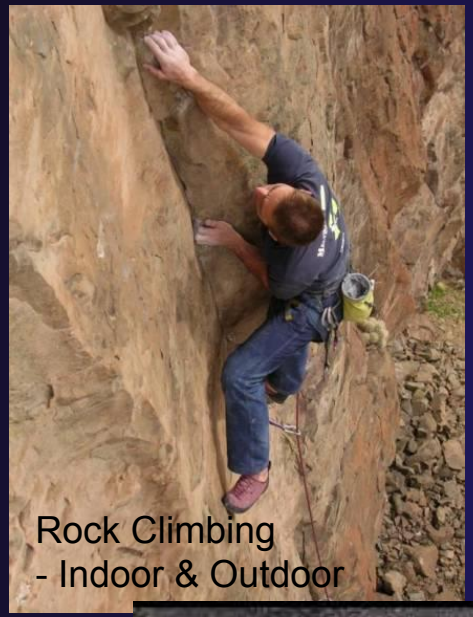

Rock Climbing
- Indoor & Outdoor

Historic Santa Fe Plaza

Mountain Biking

Rafting on the Rio Grande

Pajarito Mountain Ski Area

Santa Fe Opera

Fast Facts

32.4% Female

Veterans
6.9%

Individuals with disabilities
3.8%

#M Millennials in workforce
35.2%

Our student and postdoc pipeline is crucial for recruiting the workforce of the future

- 1,860 students and 400 postdocs were part of our workforce in FY19
- Conversion of postdocs to technical staff is our most highly utilized early-career pipeline

Percentage of total LANL population who were former students or postdocs

36%	61%	33%
All LANL employees (Reg, TRMA)	**All R&D scientists & engineers**	**Managers**

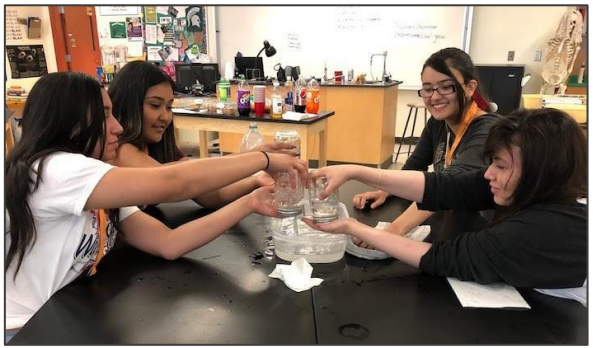

Summer Physics Camp for Young Women

Supercomputing Challenge

29% of regular/term employees have at least 1 degree from a NM college/university

41% of Los Alamos employees are native New Mexicans

An Active Community of Employee Resource Groups

American Indian Employee Resource Group

Hispanic Opportunities of Los Alamos ¡HOLA! LOS ALAMOS NATIONAL LABORATORY

Individuals with Disabilities

IWD ERG member Lani Seaman

एकता 友谊 Asian Pacific Islander ERG ความสามัคคี Chão mừng

African American Partnership Program

Atomic Women

LANL Women's Group

Women of Computing and Advocates

What We Do in Computing

- Data Science
- Cyber Security
- Software Engineering
- HPC Research
- Solving Big Data Problems
- Scientific Modeling and Simulation

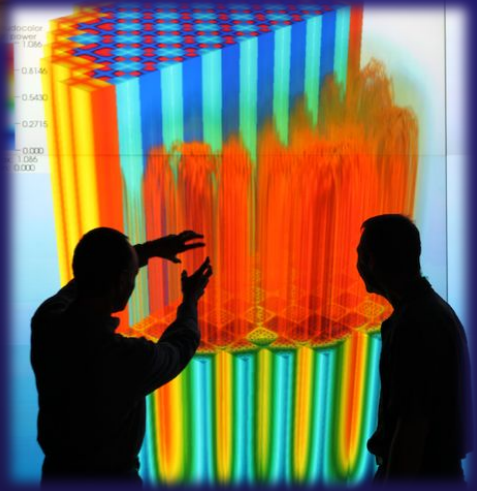

What is Scientific Computing?

- **Scientific computing is using advanced computing technology to solve complex problems**
  - Modeling and simulation of complex phenomena using supercomputers
- **National laboratories have some of the fastest supercomputers in the world**

Trinity: Los Alamos National Laboratory

- Memory Capacity: >2 PB of DDR4 DRAM
- Peak performance: >40PF
- Compute Nodes: 19,000

Example Computer Modeling Applications

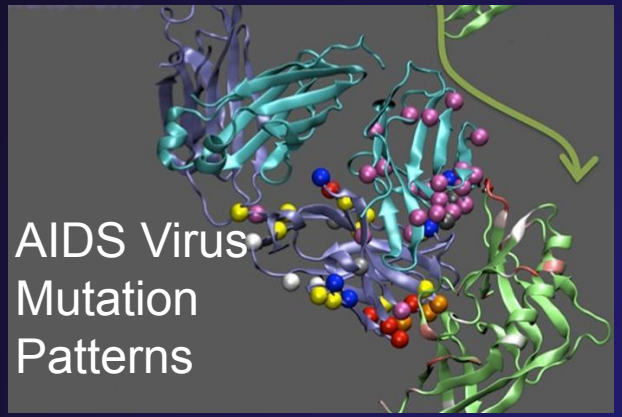

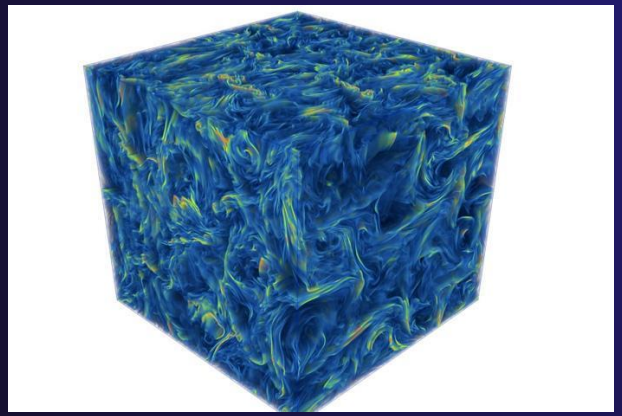

Fast Magnetic Reconnections
(sheds light on solar flares)

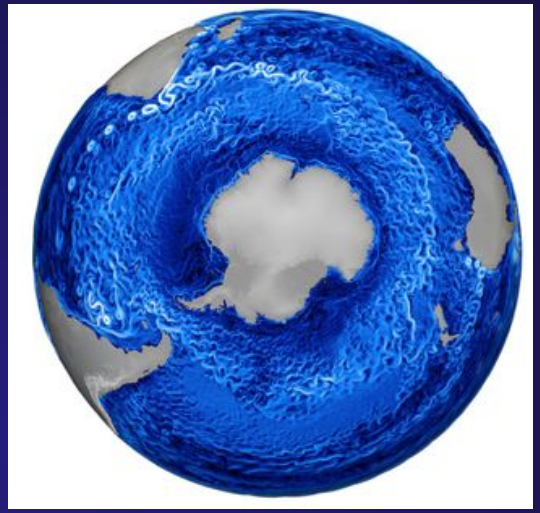

Ocean Currents
And Eddies
Around
Antarctica

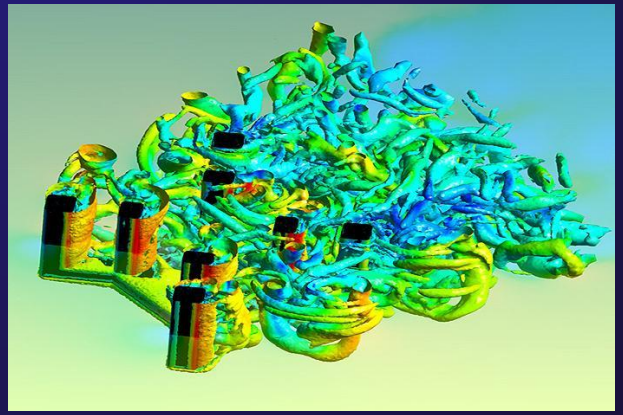

Turbulence
Around Oil
Rig Platforms

Better Climate Models? Just Add Ice

- Supercomputers provide the number-crunching muscle to model the Earth's complex global climate.
- LANL scientists used the Lab's supercomputers to couple multiple models of Antarctic ice-sheet movement to sharpen the predictive capabilities of global-climate models.
- Predictive models improve our understanding of climate change, and help us mitigate its impact.

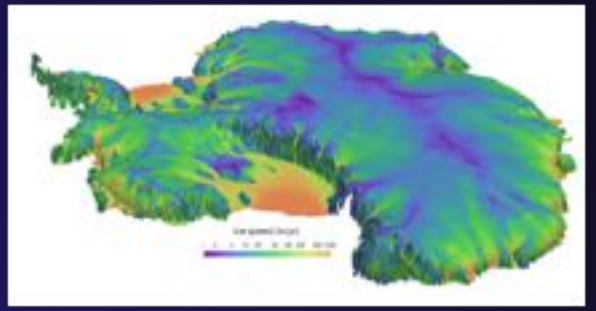

Coupling ice-sheet models with climate models, as seen here, improves the predictive abilities of both.

Slide content courtesy of Nic Lewis

More Computer Modeling Applications

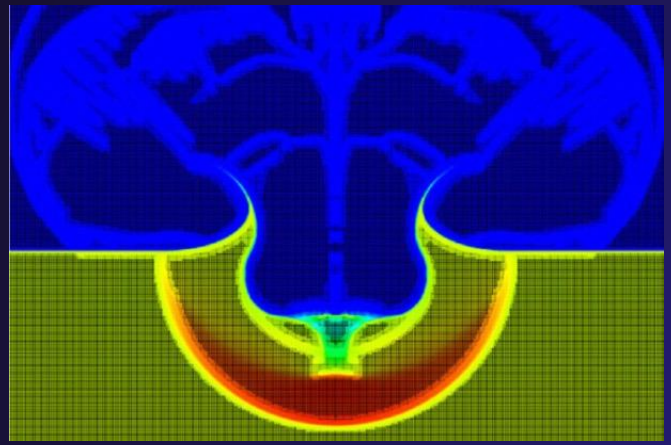

Crater simulations

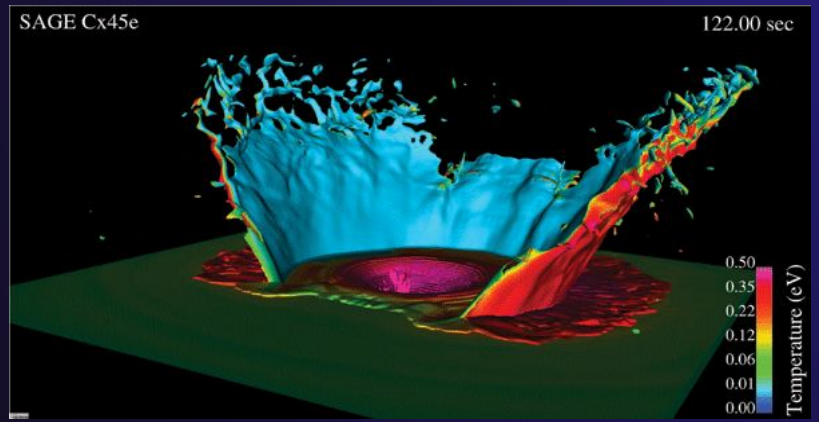

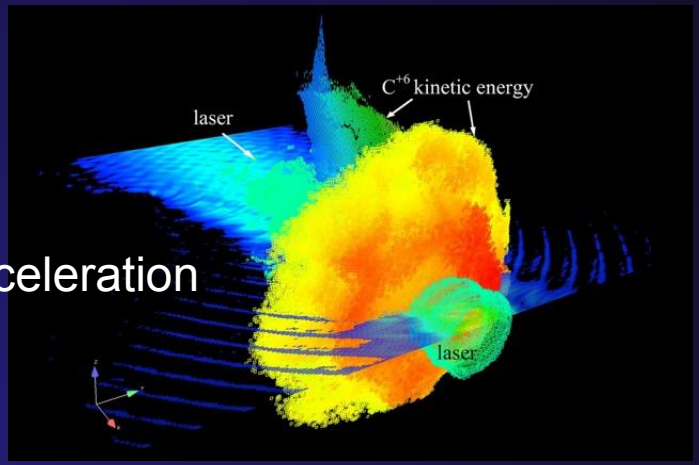

Ion Acceleration

Molecular Dimmer Switches

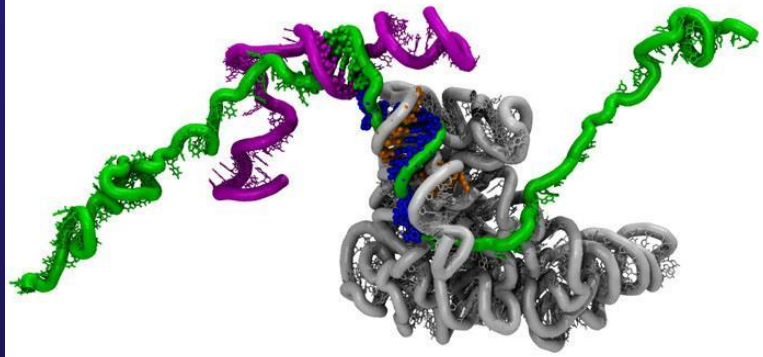

(controls cellular metabolism)

Better Science Calls for Bigger/Better Computers

Roadrunner (2007)
1st Petaflop/Accelerator Platform

Cielo (2011)
1.7 Petaflop Platform

Trinity (2015)
~20 Petaflops, 4 PB Burst Buffer

Trinity

- Cooling: Water
- Footprint: 5,200 sq ft
- Memory: 2 PB DDR4
- Burst Buffer: 3.7 PB
- Compute Nodes: 19,000

- 40 PetaFLOPs peak performance
- 80 PB parallel file storage
- 3.3 TB/s bandwidth to burst buffer
- 1.45 TB/s bandwidth to parallel storage

Designing our next Advanced Technology Platform...

Exploring and exploiting emerging technologies to improve application and workflow efficiency

How are HPC simulations done?
Killer Asteroid Example...

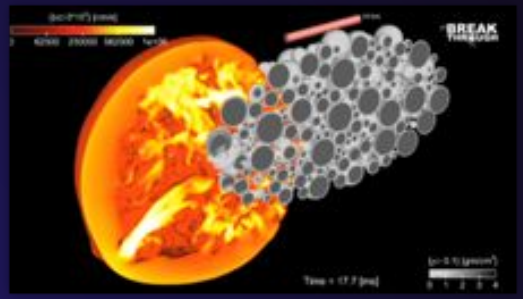

Asteroids and other Near-Earth Objects vary wildly in composition, density, and behavior. Knowing how to deflect a potential killer requires fast and accurate computer modeling on an object-by-object basis.

How are HPC simulations done?

Killer Asteroid Example – can a nuclear weapon detonate an asteroid?

- Create a computational box to do the simulation and divide into millions of cubes
- Place model of asteroid in box and assign groups of cubes to different processors.
- Allow processors to compute how contents of each cube evolve.

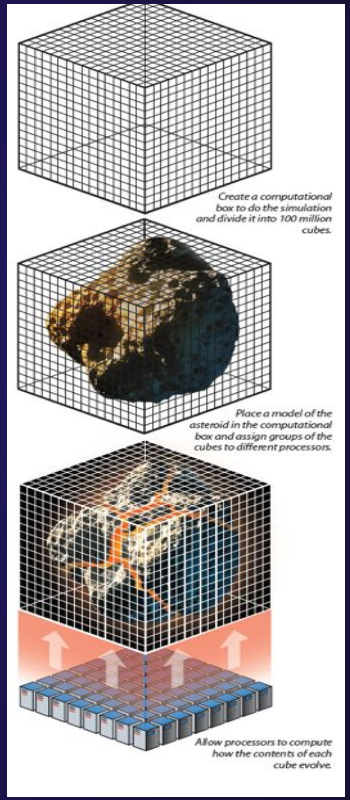

Simulations explained continued…

- **Processors simulate the event one time step at a time. When a processor computes that fragments of rock and vaporized rock in one of its assigned cubes are crossing into a neighboring cube, the processor must pass its latest data about their position, density, temperature, velocity, and so on to the processor for the neighboring cube.**

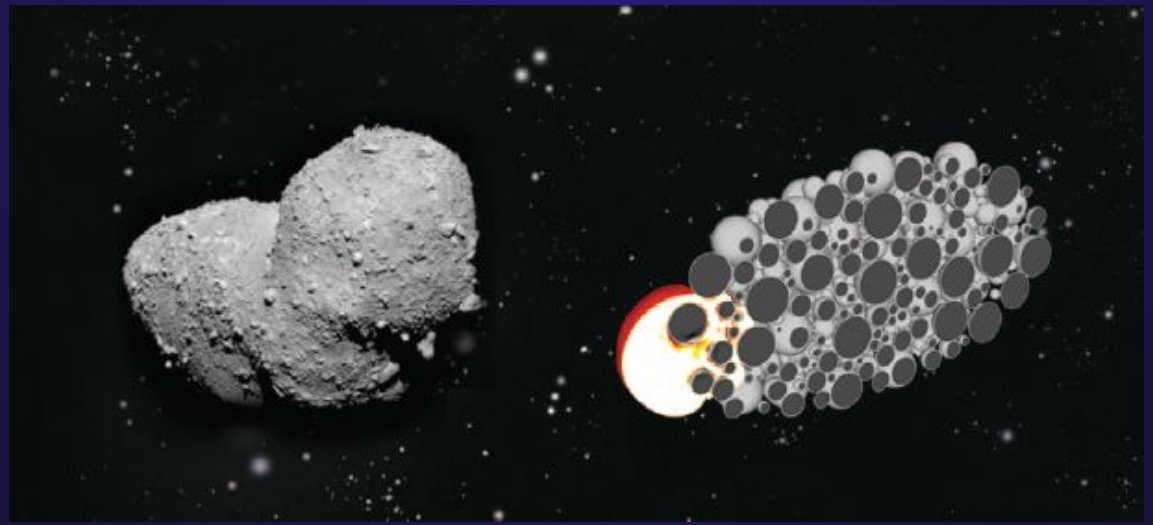

Innovation in Action:
MarFS: A Data Lake on Cloud Objects

- Provides a scalable, POSIX-friendly data lake for Enterprise, High-Performance Computing, and “big data” use
- Enables legacy systems/applications to take advantage of Cloud-based object storage
- Offers scalable data–spreads data across multiple industry storage solutions
- Provides scalable metadata–spreads metadata across multiple industry POSIX file systems

2016 R&D 100 Finalist

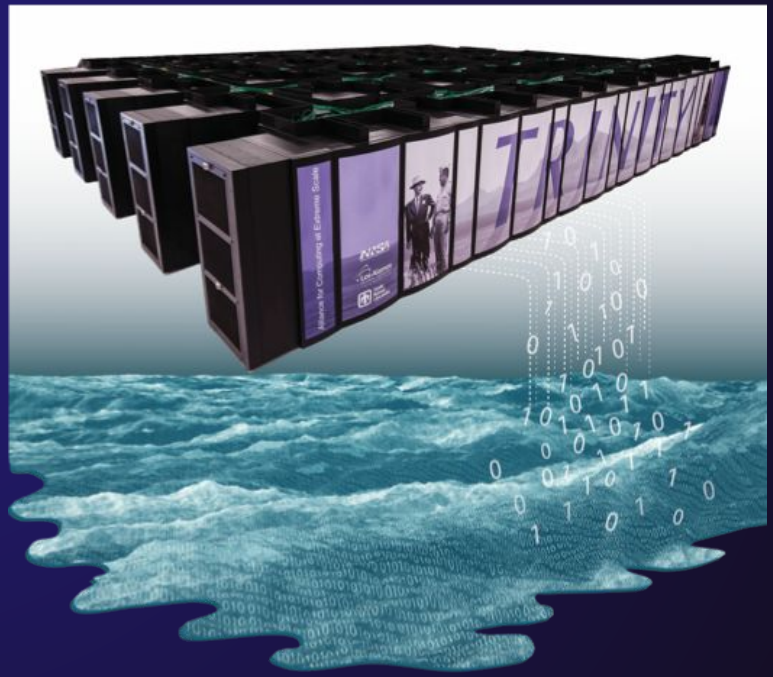

Innovation in Action:

The DeltaFS open-source distributed file system for massively parallel applications

- creates, updates, and manages extreme numbers of files,
- Alleviates the metadata bottleneck and accelerates highly selective queries.
- creates billions of files per second and does not require any additional compute resources or post-processing to create its data index.
- The performance and scalability capabilities that DeltaFS introduces are critical for storing and accessing data in the era of exascale computing.

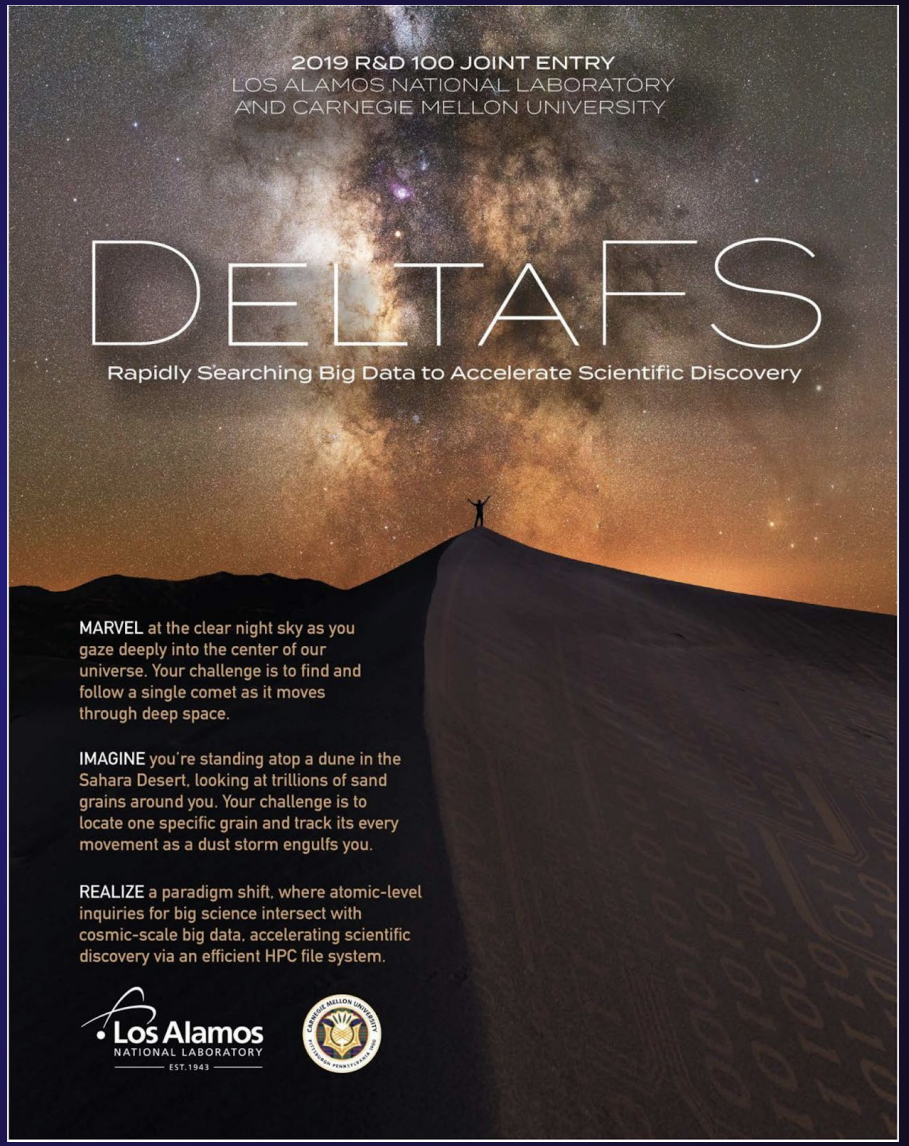

Innovation in Action:
Charliecloud

- Enables software containers - packages of custom code, software or software environments - on high performance computers.
- Achieves portability, consistency, usability and security in 1000 lines of open-source code.
- Runs on existing HPC systems with zero configuration, servers or extra processes.

Innovation in Action:
Networking Monitoring and Adaptive Routing

- IBMon/HSNmon open-source monitoring capabilities developed at LANL to allow for better real time error and performance monitoring of high speed networks (InfiniBand and Omni-Path focused)
- Dead Gateway Detection (DGD) and Neighborless Route Detection (NRD) allow for dynamic routing capabilities for supercomputers to automatically detect router failures between compute networks and Infrastructure networks

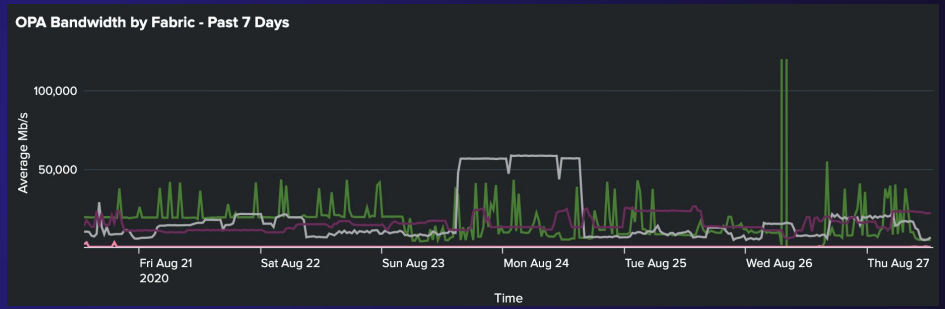

HPC@LANL YouTube Playlists

https://www.youtube.com/channel/UCJoKqqCj2DON2kFJXJrnhWg/playlists

Agenda

Trinity Supercomputer at Los Alamos National Laboratory

- LANL HPC Introduction
- **Information Security Matters**
- Network (The Nervous System - Tying it All Together)
- Data Matters: Data at Rest and Data in Motion
- Exploring Alternative and Complementary Approaches
- Hands on!

HIGH PERFORMANCE COMPUTING

Information Security Matters

The work being done in HPC demands security in both information as well as infrastructure

- Data Permissions
  - Need to know
  - Permission only given to those working on data
- Data integrity
  - Data needs to be scientifically accurate
  - Data can be collected over months of runs which can be costly to reproduce
- Open collaborative network
  - Internet accessible but still locked down
  - Work with other labs and institutions
- Dark site on restricted networks
  - Not internet accessible
  - Very secure and locked down

Agenda

Trinity Supercomputer at Los Alamos National Laboratory

- LANL HPC Introduction
- Information Security Matters
- **Network (The Nervous System - Tying it All Together)**
- Data Matters: Data at Rest and Data in Motion
- Exploring Alternative and Complementary Approaches
- Hands on!

HIGH PERFORMANCE COMPUTING

Network (The Nervous System - Tying it All Together)

- Networking is how everything is connected and how we make data and computation available to users
- Types of Networks
  - LAN (Local Area network) - Think home network
  - WAN (Wide Area network) - Think campus/town
  - Internet - No explanation there!
  - Others! Especially for HPC!
    - High Speed Networks (InfiniBand, Omni-Path, etc.)
    - Storage networks (Fibre Channel, NVMe, SAS/SATA)
  - All these networks tie in to allow us to get where we need to go!
    - HPC <-> LAN <-> WAN <-> Internet
- How is that all done?
  - This is important to know how to secure your network and understand how it works!

Network Routing and Switching

- All networking is done with switching and routing
  - Your home routers act as both a switch and a router

- **Switches** - Allow you communicate over your LAN easier and within your same subnet
- **Routers** - Allow you to leave your LAN and “route” out of your subnet via gateways (Wifi routers work the same way)

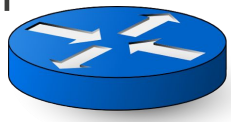

- Why is this important for security?
  - Knowing what devices can communicate locally over a switch has a much greater security impact than devices that have to route
    - If traffic stays local, that means that it won’t route or connect to external networks
  - Switch separation (Virtual LANs, or VLANs) help ensure that not everyone on the same switch can talk to each other without routing
  - HPC is vast! We need to know what is going on everywhere!

HIGH PERFORMANCE COMPUTING

Network Firewalls

- How do we protect against devices that are allowed to be “switched” or “routed” on a network?
- **Firewalls:** help enforce traffic on your network
  - Hardware firewall that everything must be “approved” through before continuing on
    - Verified source/destination of traffic before entering/leaving network
    - Access-Control Lists (ACLs)
      - Help govern traffic from an IP or mac-address perspective
  - Host based firewalls
    - Windows firewall, OSX firewall, linux iptables/firewalld
    - Can be the first or last line of defense from a network perspective
- Firewalls can also keep track of your “session”, matching multiple packets to the same connection
  - What about beyond just governing traffic?

HIGH PERFORMANCE COMPUTING

What does it all look like?

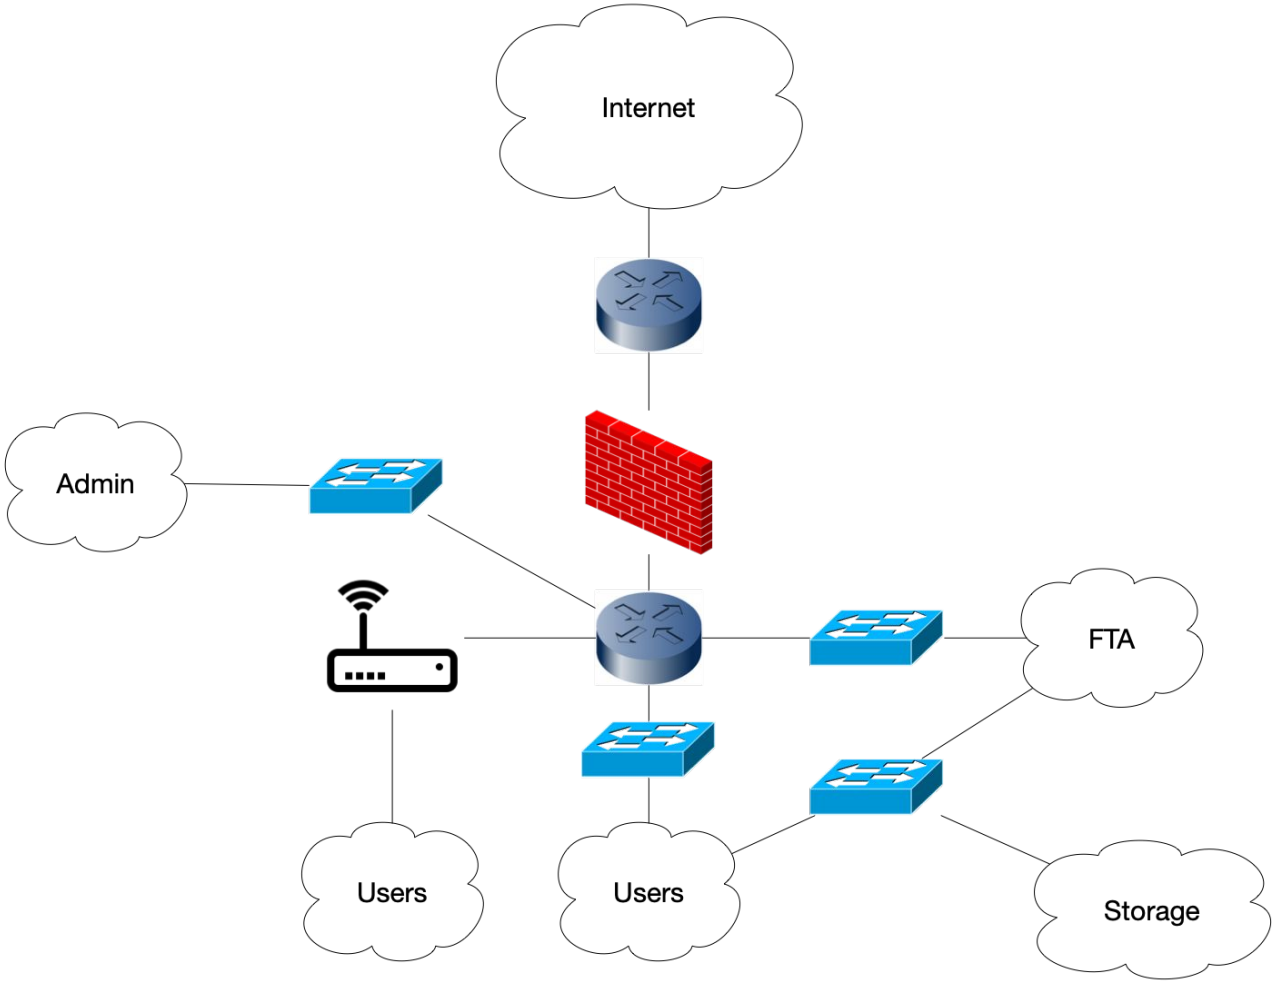

Network Intrusion, Detection and Prevention (IDP)

- What makes a good forensics analysis?
  - Logging, Logging, Logging!
  - You are only as good as your logging capabilities
- Network firewalls can do their own analysis based on traffic flows and packet info
- IDP capabilities include
  - Signature analysis (does this type of packet resemble something malicious?)
  - Filtering (allowing or blocking known sites)
  - Screening (Does it look like a particular host is poking around?)
  - Flow/packet analysis (My firewall let this traffic through, but what is it really doing?)

HIGH PERFORMANCE COMPUTING

Monitoring/Inspecting

- All network traffic can be logged any many different levels in its path
- Network tools help to provide more visibility
  - Flow analysis helps to correlate data to system logs
    - Services like sflow, Bro, security appliances, etc help to correlate all data in your environment
  - Packet inspection is ensuring the traffic coming in is actually what it's meant to be (SSH, HTTP/HTTPS, etc.)
  - Audit your traffic
- Downsides: Performance
  - HPC relies on generating lots of data from our compute clusters that users need to be able to move to multiple levels of storage tiers as well as other sites for analysis
  - How do we build a network for performance and security?

Network Trade-Offs

- Lots of network inspection and traffic hops may start to hurt your network performance
  - This matters very much in HPC where we have to move GB, TB or even PB worth of data!
- Next Generation Firewalls are designed to do a lot of inspection but for every packet that gets inspected, adds just a little bit more latency
  - (time for a packet to arrive at the destination)
  - Trying to transfer data through a firewall may hurt performance
- How can we design a network for our users to protect their data, but at the same time, ensure data movement is not restricted?

Network Best Practices

- Design your network for the best possible balance!
  - If you have a lot of users coming in and out of your network, you will want to make sure you can restrict them and monitor them
- HPC is designed with lots of backend/private networks to ensure the high-speed networks have the ability to move data and securely
  - Keep data movement away from user networks (no user interference)
  - Users interface via File Transfer Agents (FTAs) or user login nodes
- If you need to move data, provide the bandwidth and separation needed to allow for best possible performance
  - This allows for better performance, but may not have the best security
  - Ensure the data you are moving is protected in other ways!

Network Best Practices (continued)

- Firewalls should always be protecting your network (even home routers have a basic firewall to prevent access into your network)
  - Have multiple layers of firewalling (host-based) - Belt and Suspenders!
- Home routers and most industries use Network Address Translation (NAT) to go in and out of private environments
  - 192.168.1.1 may be your internal/private IP, but your home network IP is: 1.1.1.1
  - Helps to hide private IP addresses as well as give users more IPs than assigned to by Internet Service Providers (ISP)
- Separation of responsibilities
  - Let switches/routers be switching/routing
  - Let your firewall do your protecting
  - Let your services/logging be watching everything

Network Best Practices (continued)

- Segmentation
  - Separate your network (users, admins, business systems, data, etc.)
    - Can be done with VLANs, subnetting, or firewalls
    - You don't want your users in your admin or business networks!
- Demilitarized Zone (DMZ)
  - Useful for test equipment and environments that you don't want to be available to the rest of your network but maybe to the internet
  - Useful for doing "wild west" data where you may not need it to be secure
- Whitelisting/Blacklisting
  - Know what traffic you do what to allow and from where, restrict it as much as you can
  - Block known bad networks!

HPC Network Design

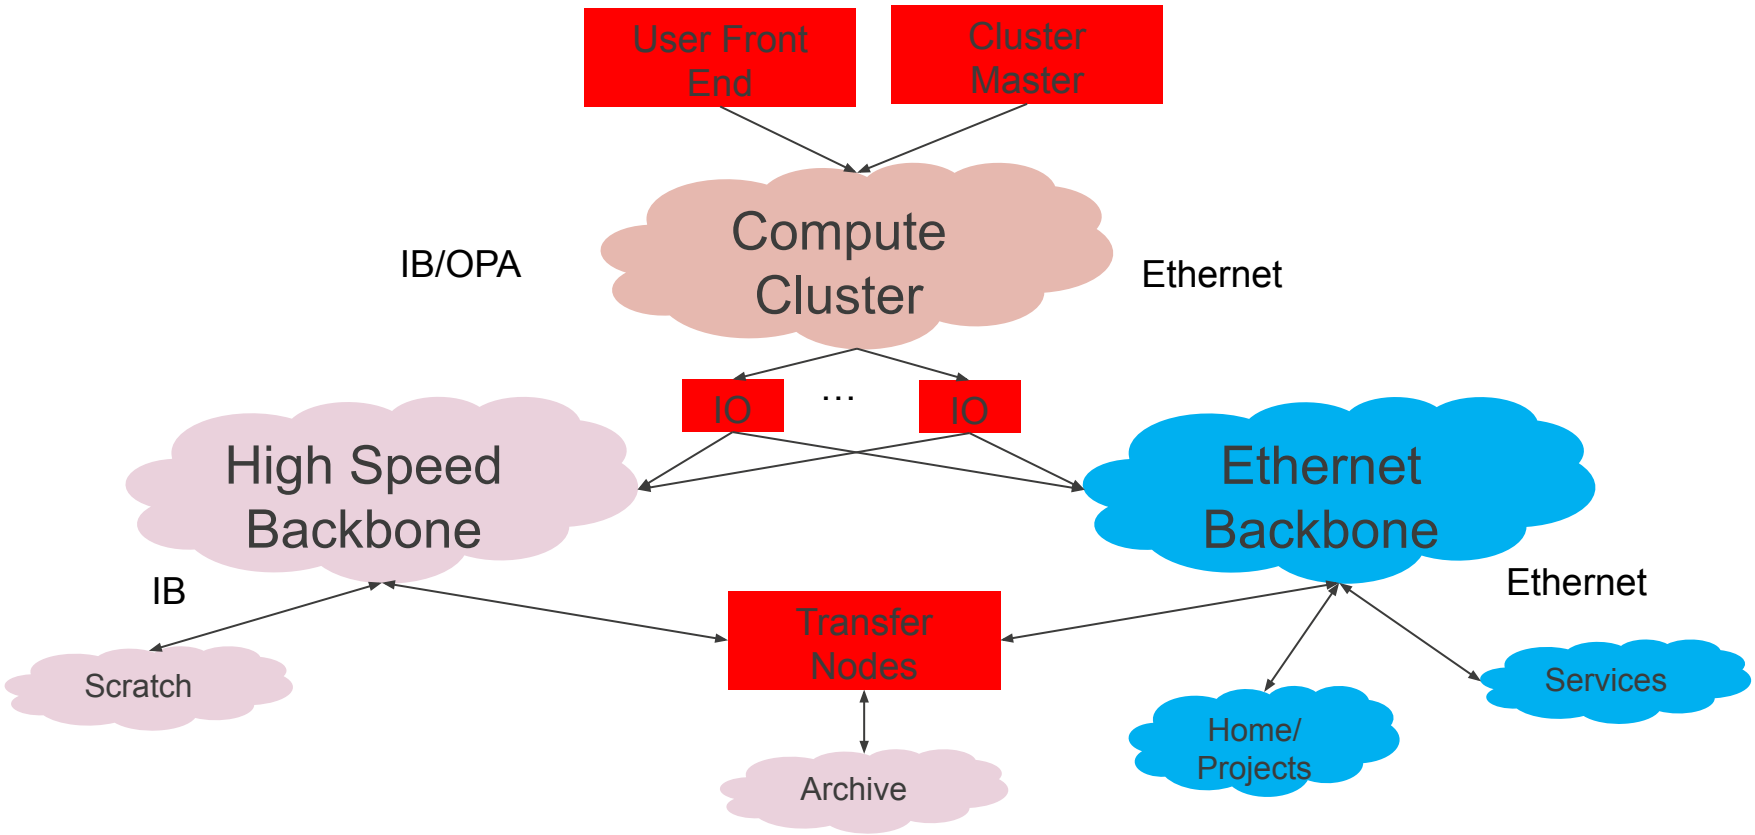

Agenda

Trinity Supercomputer at Los Alamos National Laboratory

- LANL HPC Introduction
- Information Security Matters
- Network (The Nervous System - Tying it All Together)
- **Data Matters: Data at Rest and Data in Motion**
- Exploring Alternative and Complementary Approaches
- Hands on!

Data Matters: Data at Rest and Data in Motion

Data in Motion Challenges: LANL data sets can be quite large (100TB+). Moving data efficiently and securely at that scale presents challenges with both internal and external data movement.

Real LANL external data movement example

- User runs their code on cluster at another National Lab
- They want to move it off Parallel FS at source Lab and into the Archive for safe keeping at LANL
- Data is security sensitive
- 100TB+ data sets
- Want to move it in days, not weeks

Data Matters: Data at Rest and Data in Motion (continued)

What infrastructure is needed to get data to LANL securely/efficiently?
What challenges are involved for a system architect/sysadmin?

- User vetted for transfer at remote Lab - account/authentication infrastructure
- Need beefy transfer nodes
- Need dedicated Fibre Trunk between Labs - no competing traffic
- This is sensitive data - need encryptors at each endpoint
- How do you efficiently ingest that much data at your Archive endpoint?
  - Tape doesn't do well with non-streaming workloads
  - Need enough disk to match network bandwidth
  - Need sw to parallelize transfer down to tape to mirror disk bw
- Remote computing has a cost.
  - Just the infrastructure here is $1M+ /yr
  - System Administrators! Lots of moving parts

Data Matters: Data at Rest and Data in Motion (continued)

Data at Rest Challenges: As you've learned, LANL has PBs of data, millions of dir/files across many storage tiers.

Managing this much data not easy. Users manage movement of own data between storage tiers to avoid purges. Users also need to manage their capacity usage.

This model depends on users knowing about their data

– Where did it get written?
– Does it need to be backed up? If so, did I already save a copy?
– Good naming and hierarchy

Without active management our archive accumulates too much data
Need to provide better tools for this environment.

Data Matters: Data at Rest and Data in Motion (continued)

GUFI (Grand Unified File Index)

Developed by LANL HPC - 2019 R&D 500 Award Winner!

Provide an index over all tiers of storage

- Securely allow admins (easy) and users to share the index and tools
- Reasonable update times – keep stress on source FS low if possible
- Parallel is key -- threads
- Include xattrs
- Leverage existing technology
- Keep it simple
- Command line and web interface

Data Matters: Data at Rest and Data in Motion (continued)

- Re-create source FS tree
  - Maintain ownership and permissions on the newly created tree
  - Secure – we already depend on these permissions on the source
- Use embedded DB in every dir
  - sqlite
  - This is where all file information goes
- Threads!

Please learn more here...

https://github.com/mar-file-system/GUFI

https://storageconference.us/2019/Invited/Manno.Dominic.slides.pdf

Agenda

Trinity Supercomputer at Los Alamos National Laboratory

- LANL HPC Introduction
- Information Security Matters
- Network (The Nervous System - Tying it All Together)
- Data Matters: Data at Rest and Data in Motion
- **Exploring Alternative and Complementary Approaches**
- Hands on!

HIGH PERFORMANCE COMPUTING

Exploring Alternative and Complementary Approaches

- Automated Networks
  - Software Defined Networking (SDN) for auto sensing network utilization and configurations for best performance (keeping security in mind)
- Automated firewall determinations based on packet inspections
  - Analysis tools like Bro inspect real time data and can feed firewall configuration changes
- Virtual environments - Infrastructure on Demand (IOD)
  - Cloud services that can spin up virtual environments help keep hardware costs down, but need more attention to implementation for security
- Cloud Storage
  - Pros/Cons of data storage in cloud

Complexity of Information Security at LANL

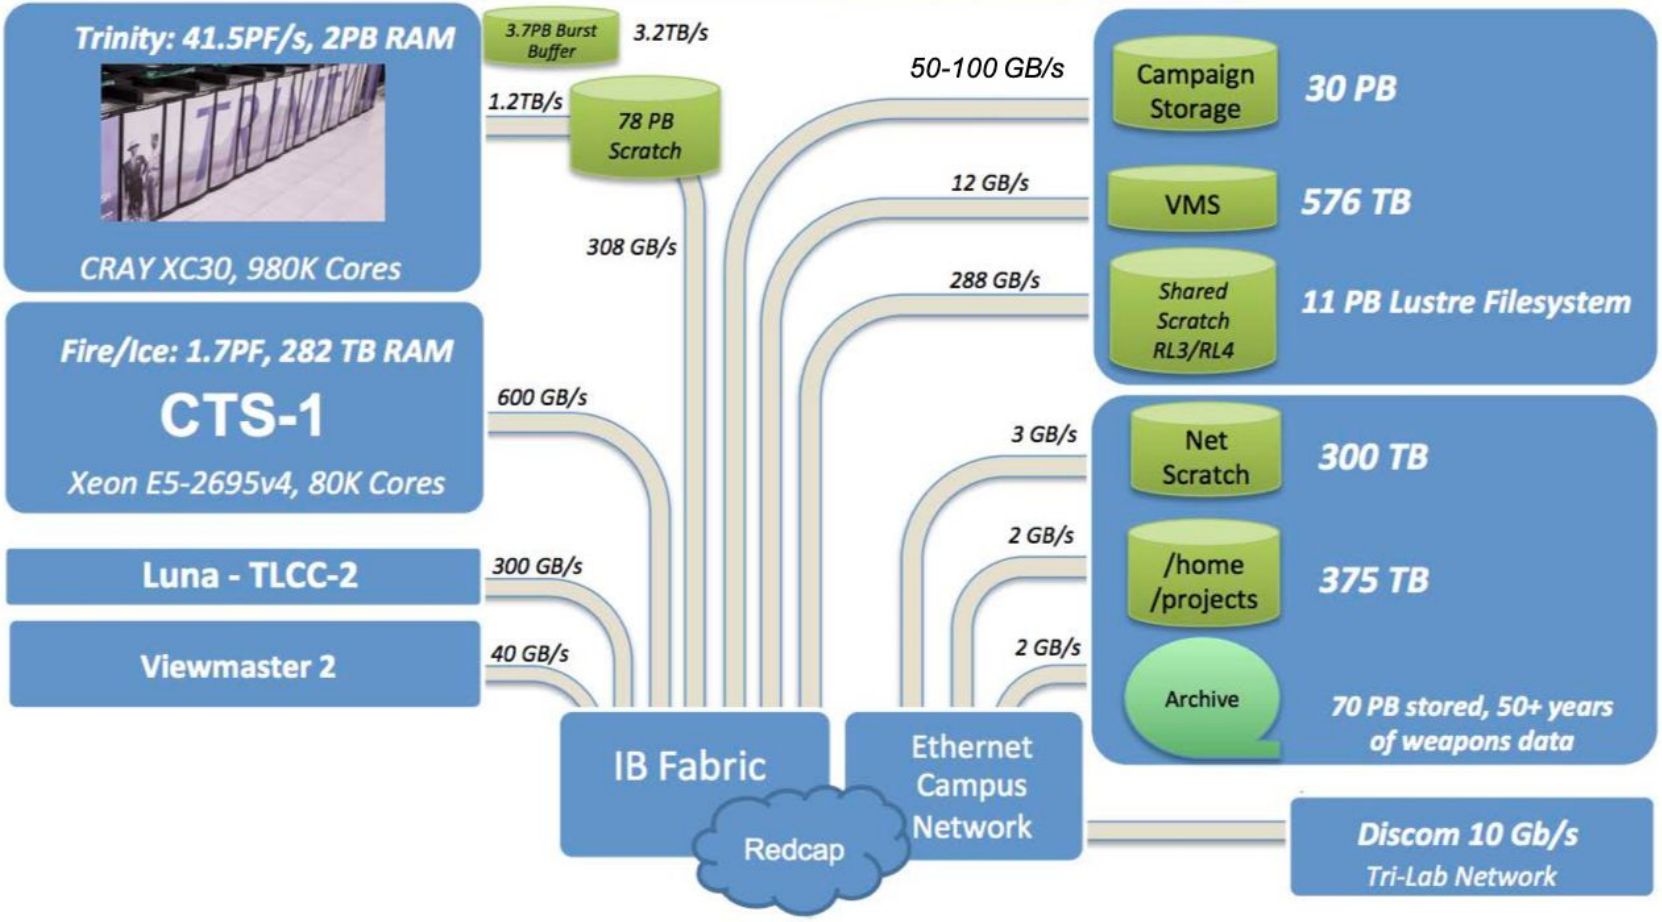

Agenda

Trinity Supercomputer at Los Alamos National Laboratory

- LANL HPC Introduction
- Information Security Matters
- Network (The Nervous System - Tying it All Together)
- Data Matters: Data at Rest and Data in Motion
- Exploring Alternative and Complementary Approaches
- **Hands on!**

HIGH PERFORMANCE COMPUTING

Now it's your turn!

What does your environment look like?

What are you or your institution doing to help protect your networks and data?

Try to find out the following on your home or institutional environments (poll afterwards):

1) What type of networking set up do I have?
2) What type of system level protections do I have?
3) What type of file system do I have on my system?
4) What am I doing to protect my data from local or external threats?

Useful commands to try on your systems to help find out that info

Windows:

Network info:

- ipconfig
- netstat -an
- route print

Firewalls:

Windows Defender Firewall

- Settings -> Windows Firewall (various ways to look at what your FW is permitting)

File System Info:

- dir
- fsutil fsinfo volumeinfo drive_letter
- fsutil fsinfo drivetype drive_letter

OSX/Linux:

Network info:

- ifconfig
- netstat --listen
- ip route

Firewalls:

IPTables

- iptables -L

firewalld

- firewall-cmd --list-all

OSX Firewall

- System Preferences -> Security/Privacy

File System Info:

- ls -l
- df -h
- du -h filename
- file filename

Poll Time!

What is your home/institution network setup?

a) NAT with private IPs and Dynamic or Static public IP
b) Public IP only
c) A or B and a firewall
d) Other

What does your computer security look like?

a) Host based firewall only
b) A & Restricted Services
c) A & B with Users/Admin accounts
d) Different combinations of A thru C
e) Other

What type of file system security is in place?

a) Singe users with file/directory permissions
b) A & multiple users/group permissions
c) A or B with multiple file systems
d) Other

HIGH PERFORMANCE COMPUTING

Thank you!

Questions?

Over 70 years at the forefront of supercomputing

Extra Slides

HPC Data Storage

- **Simulations create huge amounts of data. What are the challenges we face storing it all while providing appropriate accessibility?**

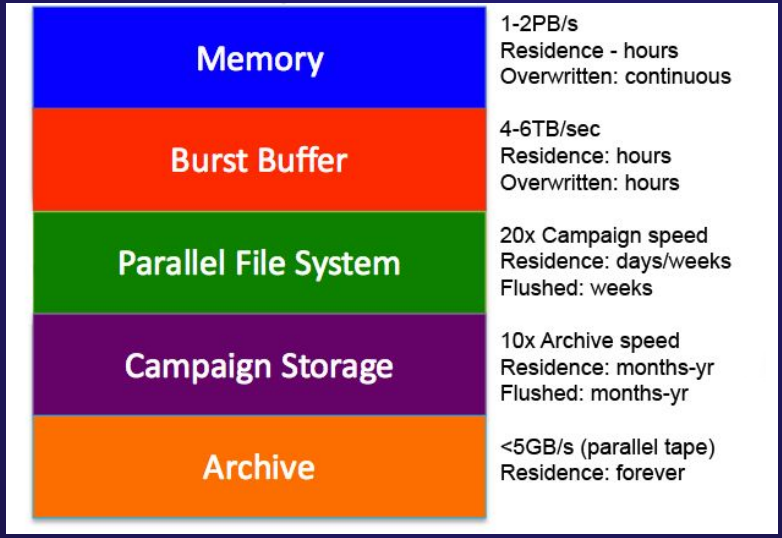

Infrastructure