# Overview of the Computer Resource Team (CRT)

Tim Fahey (Lawrence Livermore National Lab)
Ben Santos (Los Alamos National Lab)
Heidi Uphoff (Sandia National Lab)

*PSAAP III Kickoff Meeting*
*Tuesday June 30th, 2020 (Via WebEx)*
*3:00PM EDT/12:00PM PDT*

# Overview

- What is the CRT?
- What the CRT can do for you
- HPC customer support and problem tracking
- Training
- Documentation
- Communications
- Dedicated Runs
- Site visits
- Allocations
- LLNL compute resources
- LANL compute resources
- Sandia compute resources
- Getting accounts - SARAPE

# What Is The CRT?

- The **Computer Resource Team (CRT)** is the component of the PSAAP III program that connects center researchers with the High Performance Computing (HPC) resources required to perform their work.

- The CRT is comprised of representatives from each NNSA Lab who are familiar with their Lab's computing resources, personnel and policies. The following individuals serve on the CRT:
  - Tim Fahey, LLNL *fahey2@llnl.gov*
  - Ben Santos, LANL *bsantos@lanl.gov*
  - Heidi Uphoff SNL *hauphof@sandia.gov*

- ***Our primary purpose is to provide assistance and guidance in all aspects related to the use of HPC resources located at LANL, LLNL, and Sandia.***

# What the CRT Can Do For You?



- Assist with the establishment and use of computer accounts
- Assist with accessing compute resources
- Provide essential HPC user documentation
- Provide technical support and referral to in-depth consulting
- Conduct periodic telecons to keep PSAAP users up-to-date with account, access, policy, scheduling and technical issues, and to address issues with HPC platform usage
- Interface with other individuals and groups within the Labs, such as management, networking, system administration, storage, customer support, etc., to facilitate the effective support of PSAAP users
- Track and facilitate the resolution of problems reported to each Labs' customer support “hotline”
- Provide training opportunities

# What the CRT Can Do For You? (cont)



- Collect and distribute monthly machine usage statistics
- Schedule and support special/dedicated runs
- Maintain a balance of machine usage between the PSAAP centers (if needed)
- Visit PSAAP centers to discuss HPC resources, user issues and to offer technical consultation and/or training
- Showcase PSAAP research in the NNSA/ASC or DOE research exhibit booth at the annual SC conference

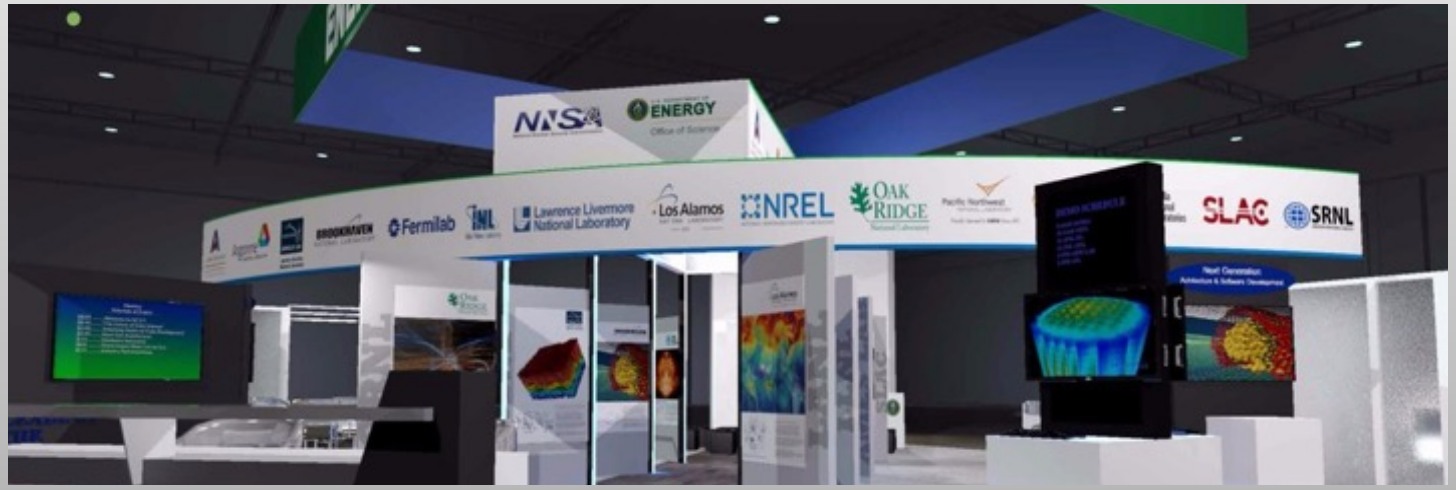

# HPC Customer Support and Problem Tracking

- All three labs offer HPC customer support via phone and email:
  - LLNL: Livermore Computing Hotline
  - LANL: ICN Consulting Office
  - Sandia: HPC OneStop Service Desk
- Includes support for user accounting issues and for technical assistance.
- Problems and questions are tracked via a customer support database application (varies with each Lab).
- Most problems/questions are handled directly by the customer support staff on duty.
- More in-depth issues are typically referred to local subject experts.
- The labs also coordinate with hardware and software vendors for issues that require outside support.
- CRT reps coordinate routinely with each other on Tri-lab user issues.

# HPC Training

- Training is important – especially for new users
  - Online tutorials are available
  - Workshops conducted at the Labs are open to PSAAP center users
  - The CRT can deliver workshops/training at your center or virtually
  - The CRT can also assist with topic specific, customized workshops if that is of interest
- Topics include:
  - Getting started information
  - Compilers
  - Performance tools
  - Debuggers
  - Parallel programming (MPI, OpenMP, Pthreads)
  - Batch schedulers
  - Machine architectures
  - Visualization tools
  - …

# User Documentation

- Most of what users need to know is available online via web pages hosted by each of the Labs. Recommended starting points:
  - LLNL
    - https://hpc.llnl.gov/user-portal
    - https://hpc.llnl.gov/training/tutorials
  - LANL
    - https://www.lanl.gov/org/ddste/aldsc/hpc/index.php
    - https://hpc.lanl.gov (LANL Cryptocard)
    - https://hpcinfo.lanl.gov (Dashboards)
  - Sandia
    - https://computing.sandia.gov
    - http://hpc.sandia.gov (open)
  - Note: most LLNL web pages are open – no authentication required. Most Sandia / LANL pages require authentication

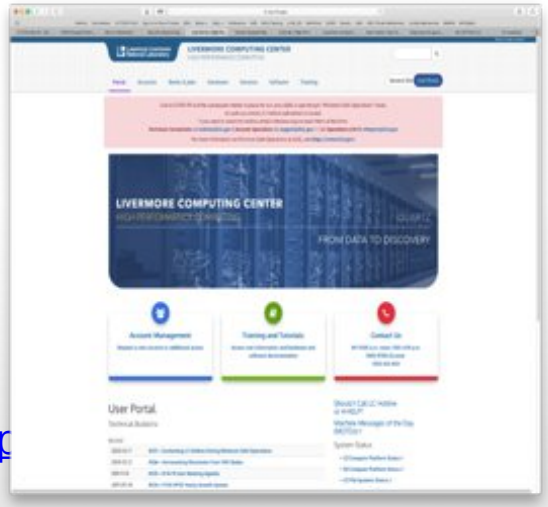



- Quick Start Guide for new PSAAP users:
  - https://asc.llnl.gov/content/assets/docs/Alliance.Quickguide.pdf

# Communications

- Monthly telecons
  - Forum for discussion/questions on user topics such as accounts, access, technical issues, machine schedules, etc.
  - Open to all PSAAP center users
  - Highly recommended to have a Point-of-Contact (POC) person with some technical/computing experience from each Center attend. We'll be in touch with each Center's PI to find out who this might be.
  - Day and time for these monthly WebEx's will be set at a future date.
- Email list - psaap-crt-llnl@llnl.gov
  - Telecon minutes and machine usage statistics are distributed via this list to all PSAAP Center PIs & POCs, ASC HQ and various staff & managers within the Labs
  - Let us know if you want anyone else at your Center added to our list - initially it includes only your PI and POC

# Communications



- Usage stats
  - Collected by each Lab's computing center and distributed with the telecon minutes
  - Present both aggregate and detailed usage (down to the user level)

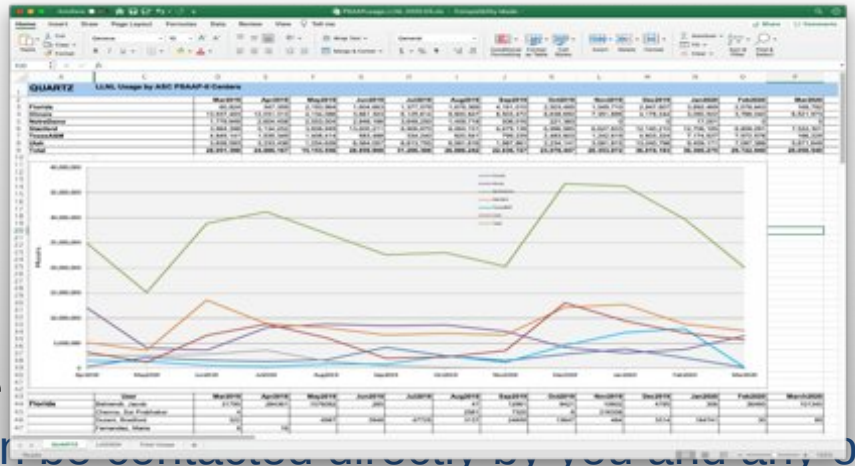

- Email & phone
  - The CRT can be contacted directly by you and any of your Center's users:
    - Tim Fahey (LLNL) *fahey2@llnl.gov* 925-422-4228
    - Ben Santos (LANL) *bsantos@lanl.gov* *505-665-6153*
    - Heidi Uphoff (Sandia) *hauphof@sandia.gov* 505-844-6119

# Dedicated Runs (DATs)



- Normally, users submit their jobs to a batch system using a standard Tri-lab SLURM or LSF batch scheduler interface:
  - Jobs are then queued to wait their turn for execution.
  - There are limits on the number of nodes and number of hours that a job can use.

- Dedicated application time - DAT: can be requested by any PSAAP user:
  - Overrides normal node limits - up to the full machine
  - Overrides normal time limits
  - Typically conducted on weekends at LLNL, LANL date flexible but prefer 2-week lead time
- How to request a DAT:
  - LLNL: https://hpc.llnl.gov/accounts/forms/asc-dat
  - LANL: https://hpc.lanl.gov/dat_request

# Site Visits



- Historically the CRT has been available to visit PSAAP centers:
  - Visits (2-4 hrs) by the CRT and other Lab staff
  - Focus is on the center's *users* of HPC computing resources
  - Updates on architectures, policies, future plans at the Labs
  - Forum for discussing user issues, problems, questions
  - Not to be confused with TST visits or reviews
- Technical "training" sessions were offered if desired
- In our current environment is is highly likely these will be virtual. The CRT will work with each center to schedule these sessions.

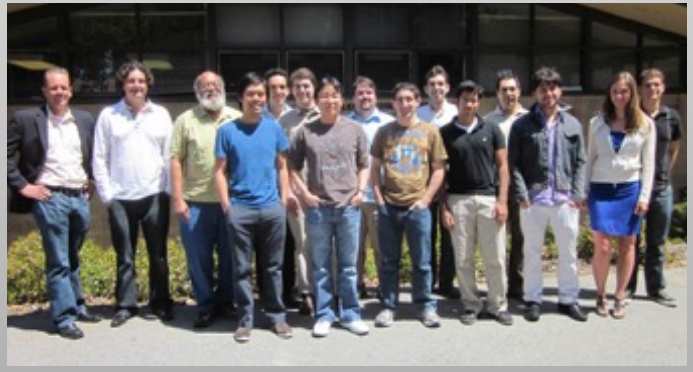

# Allocations Production Machines

§ **PSAAP centers have allocations on eight ASC funded, Tri-lab production machines:**

- **PSAAP allocations vary by machine; shared at the group level**
- **Able to use more than allocation if there are available cycles**

| Machine | Lab | Total available core-hr/mo | PSAAP allocation | PSAAP total core-hr/mo |
|---|---|---|---|---|
| QUARTZ | LLNL | 78,209,280 | 31.25% | 24,440,400 |
| LASSEN* | LLNL | 25,310,560 | 10.0% | 2,531,056 |
| SOLO | Sandia | 9,828,720 | 1.0% | 98,287 |
| ECLIPSE | Sandia | 39,104,640 | 1.0% | 391,046 |
| ATTAWAY | Sandia | 39,104,640 | 1.0% | 391,046 |
| GRIZZLY | LANL | 39,157,200 | 1.0% | 391,572 |
| SNOW | LANL | 9,671,040 | 1.0% | 96,710 |
| CAPULIN | LANL | 7,154,000 | 1.0% | 71,540 |

*** Lassen allocation is subject to change**

NNSA National Nuclear Security Administration · ASC · Los Alamos National Laboratory · Sandia National Laboratories · Lawrence Livermore National Laboratory

# LLNL PSAAP Compute Resources

- **Quartz**
  - Intel Xeon E5-2695v4 (Broadwell) architecture;  2.1 GHz
  - 3018 nodes with 36 cores each
  - 108648 total cores
  - 128 GB memory per node; 344,064 TB total memory
  - Omni-Path interconnect; ~100 Gb/sec pt2pt MPI
  - 3.23 petaflops peak
  - Details at https://hpc.llnl.gov/hardware/platforms/Quartz

# LLNL PSAAP Compute Resources



- **Lassen**
  - IBM Power 9 architecture at 3.5 GHz
  - 792 compute nodes with 44 cores each; 34,848 total cores
  - 4 NVIDIA V100 GPU's per node; 3168 total GPU's
  - 256 GB CPU-memory and 64 GB GPU-memory per node
  - Infiniband EDR interconnect
  - 23 petaflops peak
  - Details at https://hpc.llnl.gov/hardware/platforms/lassen

# LLNL PSAAP Compute Resources



- **HPSS Archival Storage**
  - Mounted from all unclassified production clusters
  - Users automatically get an HPSS account with their machine accounts
  - 300TB quota for FY2020
  - 154 PB total archive storage
  - 6 PB disk cache

- **Visualization Cluster – Pascal**
  - Intel Xeon E5-2695 v4 (Broadwell)
  - 164 nodes, 36 cores/node
  - 2 NVIDIA Tesla P100 GPUs/node
  - 256GB memory/node
  - 16GB GPU global memory/node
  - Available upon request

# LANL PSAAP Compute Resources

Ben Santos

# LANL PSAAP Compute Resources

- **Grizzly**
  - Intel Xeon Broadwell E5-2695 2.1 GHz
  - 1490 nodes with 36 cores each; 53640 total cores
  - 128 GB memory per node; 191 TB total memory
  - Intel OmniPath interconnect
  - 1806 teraflops peak
  - 23.5 PB Lustre (global scratch)
  - https://hpc.lanl.gov/grizzly_home

# LANL PSAAP Compute Resources

- **Snow**
  - Intel Xeon Broadwell E5-2695 2.1 GHz
  - 368 nodes with 36 cores each; 13248 total cores
  - 128 GB memory per node; 47.1 TB total memory
  - Intel OmniPath interconnect
  - 445 teraflops peak
  - 23.5 PB Lustre (global scratch)
  - https://hpc.lanl.gov/snow_home

# LANL PSAAP Compute Resources



- **Capulin**
  - Cray XC50 Cavium ThunderX2 ARM
  - 175 nodes with two TX2 56 cores per node; 9800 total cores
  - 256 GB memory per node; 45 TB total memory
  - Cray Aries interconnect
  - 196 teraflops peak
  - 23.5 PB Lustre (global scratch)
  - https://hpc.lanl.gov/capulin_home

# LANL PSAAP Compute Resources



- **Turquoise Archival Storage**
  - Accessible from ar-tn.lanl.gov via wtrw.lanl.gov
  - Mounts the scratch spaces, internet staging and archive
  - http://hpc.lanl.gov/turquoise_archive

- **Fast File Transfer Service**
  - Data Transfer Nodes (DTN)
    - Globus supported endpoint (https://hpc.lanl.gov/globus_intro)
    - Utilize scp and parallel bbcp
  - Uses a 2-hop method (temporary internet staging storage)
  - https://hpc.lanl.gov/dtn

# Sandia PSAAP Compute Resources

Heidi Uphoff

# Sandia PSAAP Compute Resources




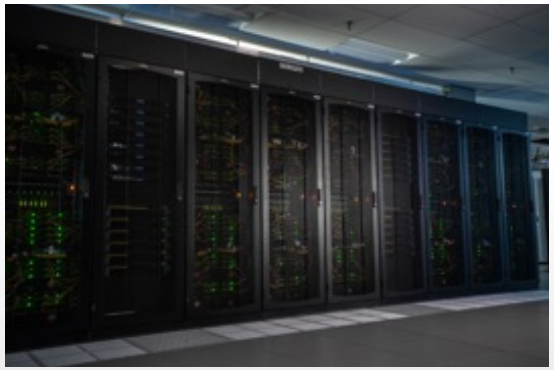


- **ECLIPSE**
  - Intel Broadwell E5-2695 v4;  2.1 GHz
  - 1,488 nodes with 36 cores each; 53,568 total cores
  - 128 GB memory per node
  - Intel Omni-Path high speed interconnect / Mellanox ConnectX4
  - 1.800 petaflops peak
  - In SARAPE, request "SRN Capacity Clusters“(US citizenship required)
  - http://hpc.sandia.gov/HPC%20Production%20Clusters/index.html

# Sandia PSAAP Compute Resources



- **ATTAWAY**
  - Intel Xeon Gold 6140;  2.3 GHz
  - 1488 nodes with 36 cores each; 53,568 total cores
  - 192 GB memory per node
  - Intel Omni-Path high speed interconnect / Mellanox ConnectX4
  - 3.943 petaflops peak
  - In SARAPE, request "SRN Capacity Clusters“(US citizenship required)
  - http://hpc.sandia.gov/HPC%20Production%20Clusters/index.html

# Sandia PSAAP Compute Resources

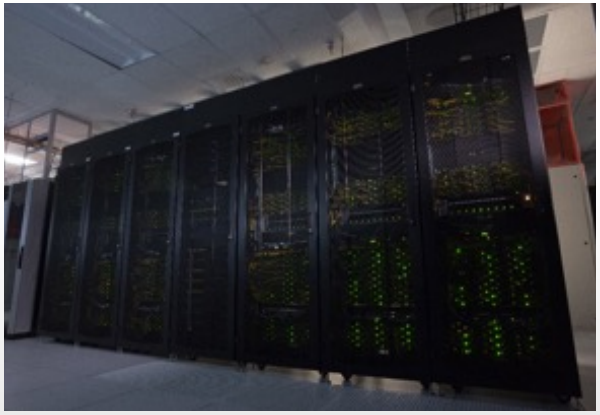

- **SOLO**
  - Intel Broadwell E5-2695 v4;  2.1 GHz
  - 374 nodes with 36 cores each; 13,464 total cores
  - 128 GB memory per node
  - Intel Omni-Path high speed interconnect / Mellanox ConnectX4
  - 460 teraflops peak
  - In SARAPE, request “ECN Capacity Clusters“
  - http://hpc.sandia.gov/HPC%20Production%20Clusters/index.html

# Sandia PSAAP Compute Resources



- **Sandia Mass Storage System (SMSS)**
  - High Performance Storage System (HPSS) provides high-end near-line storage for HPC systems
  - In SARAPE, request “Restricted Sandia Mass Storage System (RSMSS)“ for the SRN network and “ECN Capacity Clusters” for the ECN and OHPC network.

- **LYNX**
  - Sandia’s HPC file transfer agent nodes
  - Mounts the same parallel file systems available on the compute clusters
  - 10 gigabit and InfiniBand (IB) technology to optimize file transfer performance to/from Sandia SMSS
  - High performance data transfer tools HSI, HTAR, and PFTP
  - In SARAPE, request "SRN Capacity Clusters"

# Sandia PSAAP Compute Resources



- **Testbed systems**
  - Request accounts through SARAPE (sarape.sandia.gov)
  - Research systems from several vendors
  - Some require NDAs
  - Not for production computing
  - Systems come and go, lifetime usually a year or less
  - Users cannot expect mature hardware/software
  - Support available via Sandia HPC OneStop (e.g. email to <*machine-name*>-help@sandia.gov
  - In-depth support typically provided by research teams

# Sandia PSAAP Compute Resources

## Tri-lab Advanced System Technology Testbeds - A Co-Design Tool

- Not for production computing cycles - but can be provided to Test Pilot users
- Both hardware and software are intended to be highly dynamic
- Closer to prototypes and technology development drivers
- Multiple nodes available but more important to explore a diverse set of architectural alternatives, than push large scale

- Available for PSAAP access, via SARAPE.

| | Mayer | White | Blake | Morgan | Stria |
|---|---|---|---|---|---|
| **CPU** | 4 computes B0 ThunderX2 (beta)<br>43 computes A1 ThunderX2 | Dual IBM Power 8, 10 cores | Dual-Socket Intel Xeon Platinum | Five dual socket Intel IvyBridge with Xeon Phi co-processor and four Intel 32-core Haswell nodes | 2.0 GHz Arm Cavium Thunder-X |
| **Accelerator** | None | NVIDIA K40 2 per node | None | Intel Xeon Phi Co-processor (codenamed Knights Corner) 2 per node | None |
| **Nodes** | 47 | 9 | 40 | 9 | 266 |
| **Interconnect** | Mellanox EDR Infiniband with SocketDirect | Mellanox EDR IB | Intel OmniPath | Mellanox Quad Data Rate InfiniBand | Infiniband |
| **Other** | n/a | Technology on the path to anticipated CORAL systems | Each processor core has dual AVX512 vector processing units that are FMA capable. | Hetero testbed on restricted network<br><br>* On SRN network (US citizenship required) | * On SRN network (US citizenship required) |

http://www.sandia.gov/asc/computational_systems/HAAPS.html

# Computer Accounts - SARAPE

- Each Lab has its own policies, forms and procedures, however there is a single-entry portal (**sarape.sandia.gov**) for requesting an account at any of the 3 Labs.

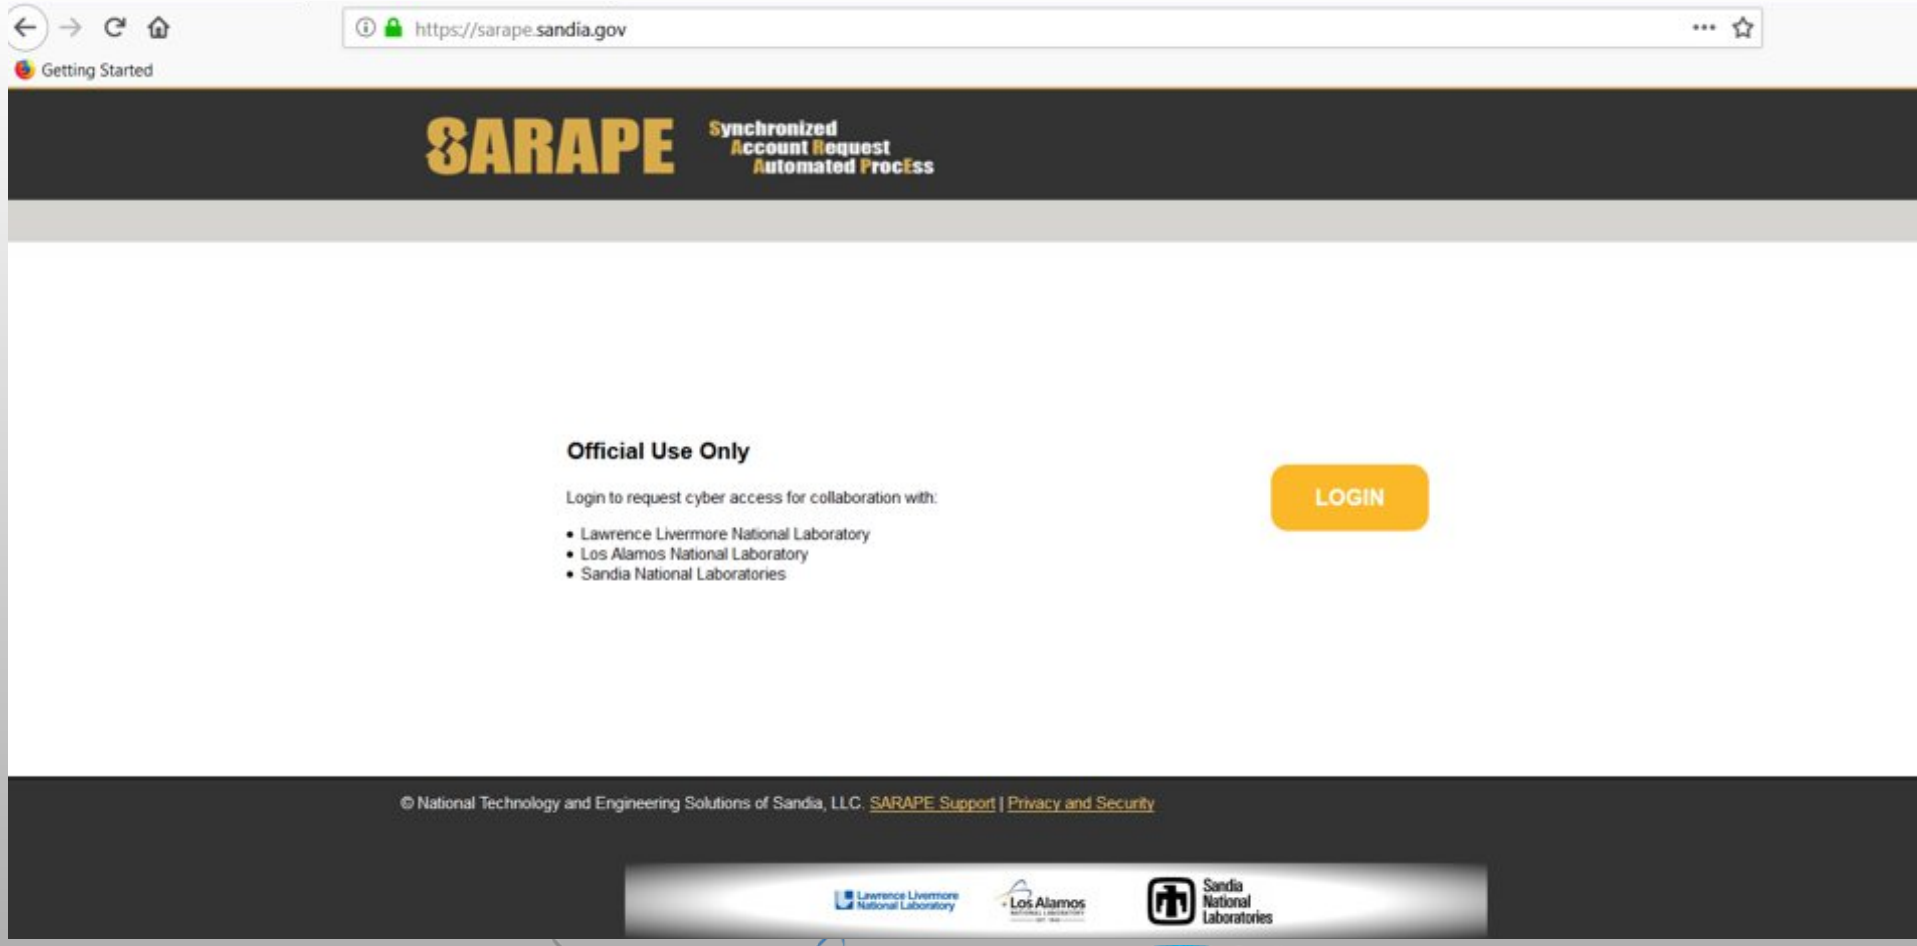

# Computer Accounts - SARAPE



- Centers need at least one *account authorizer* (also called a GPA). This can be a PI, POC and/or a trustworthy, knowledgeable designee.
- Account authorizers are responsible for overseeing and approving (via SARAPE) the accounts for all of their Center's users. Here's who we have as of today:
  - Colorado – Richard Regueiro
  - Illinois - Mike Anderson
  - Stanford – Corinne Beck, Vi Nguyen
  - Texas – Karl Schulz
  - Buffalo – James Chen
  - MIT – Jean Sofronas
  - Maryland – Johan Larsson
  - New Mexico – Tracy Wenzl
  - Oregon State - Kyle Niemeyer
- Having a “backup” authorizer is important if the primary authorizer is often not available.

# Computer Accounts - SARAPE



- After a SARAPE account request is approved by the Center's account authorizer, it is routed to the appropriate Lab for review, collection of additional information and approval.
- Typically takes 1-2 weeks for US citizens
- Account processing for non-US citizens requires additional time and “paperwork” - allow 30-90 days (plan ahead!).
- Lilia Martinez (Sandia: lmartin@sandia.gov 505-845-7967 ) manages the sarape.sandia.gov user accounts portal and has already contacted the PIs of each PSAAP II Center to get the account request process going.
- Questions? Contact Lilia Martinez or one of the CRT representatives.

# Questions?

# Backup