

# ADAC Federated Testbed

## Creating a Blueprint for Portable Ecosystems

Sadaf Alam, Jeffrey Vetter, Mark Klein,  
Maxime Martinasso, ExCL team @ ORNL, ...

ADAC Workshop  
February 15, 2018



January, 2016



June, 2016



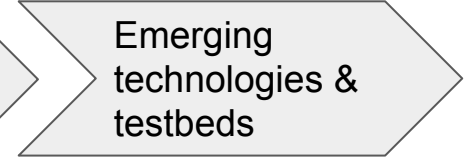
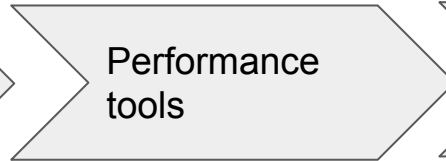
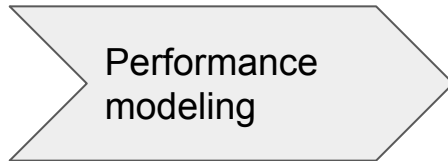
January, 2017



July, 2017

# Motivation and Use Cases

## ADAC 1-4 workshops





ExCL and  
CSCS  
experimental  
testbeds

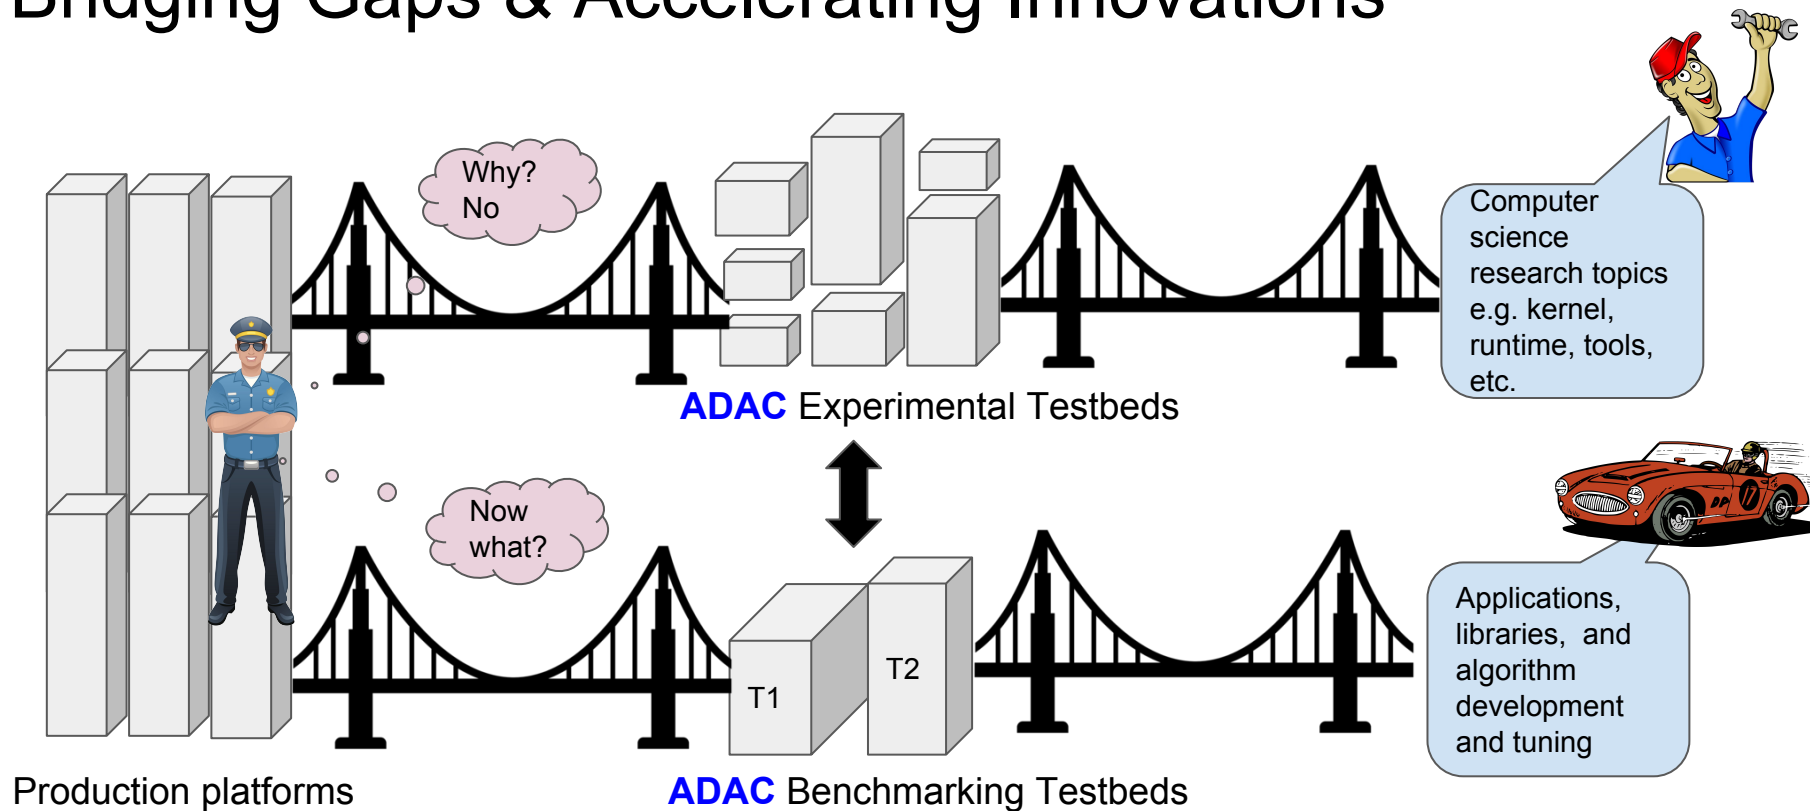


TokyoTech,  
ORNL LCCF and  
CSCS evaluation  
systems



*"GEOFFREY'S KIT CAR CAN GO FROM NOUGHT  
TO UPSIDE DOWN IN UNDER NINE SECONDS"*

# Bridging Gaps & Accelerating Innovations



# Summary: Motivation & Use Cases

- Enable evaluation of full ecosystem for emerging technologies
  - Including experiments requiring privileged access
  - Hardware level testing
  - Memory and storage hierarchies
- Share testbed resources
  - Inventory of resources
  - Federated access
  - Coordinating installation of new hardware
- Share benchmarking and evaluation results
  - Portfolio of results
  - Collection of tools with support matrix

The screenshot shows a web browser window with the URL <https://3.basecamp.com/3618876/projects/4785095>. The page title is "ADAC Emerging Technologies Working Group". Below the title is a row of avatars for team members: AN, BM, DE, SA, and MM, followed by a link "Add/remove people...". The page is divided into three main sections: "To-dos", "Message Board", and "Schedule".

**To-dos** 2/7

**Action Items**

- ☐ list of testbeds at each site  
Oct 6, 2017 Jeffrey V. Akihiro N. Maxime M. Scott A.
- ☐ strategy for engaging vendors on testbed systems  
Oct 3, 2017 Jeffrey V. sadaf a.
- ☐ list mechanisms for providing access to testbeds (at each site)  
Oct 2, 2017 Jeffrey V.

**Message Board**

- BM Info for ADAC Exec meeting at SC
- BM Demonstration for Tokyo Meeting
- SA Define capabilities for a testbed system management
- SA Testbed Resources  
ORNL NCCS Resources System

**Schedule**

Feb 15 - Feb 16  
Tokyo meeting - tentative date

Monthly calls + topic related discussions

# Inventory (Work in Progress)

# Experimental Computing Laboratory (ExCL)

## Status update Dec 2017

- **Growing pains— Improved infrastructure**

- ExCL expanding into Annex
- 40Gb optical links connecting Annex to JICS
- New fileserver, utility servers
- Bare metal provisioning

- **Upcoming Resources**

- FY18Q2
  - IBM TrueNorth
  - D-Wave (managed access to their cloud)
  - Zynq SoC
  - Atos Quantum Learning Machine



By DARPA SyNAPSE - [http://www.darpa.mil/news\\_events/telecom/2014/05/07.aspx](http://www.darpa.mil/news_events/telecom/2014/05/07.aspx), Public Domain, <https://commons.wikimedia.org/w/index.php?curid=3465417>

- **Current**

- Emu Chick system
- Arria 10 FPGA
- Stratix V FPGA
- Pascal P100
- Fusion IO



# Inventory: CSCS Experimental Testbed

- Existing infrastructure
  - Traditional with
  - Multiple generation Xeon based systems with GPUs
  - InfiniBand and OPA interconnect and HCAs
  - Power 8
  - ARM8 nodes
- Upcoming
  - AMD CPUs and GPUs
  - Intel FPGAs
  - SSD storage such as
  - To be announced ...

# Inventory: Evaluation and Benchmarking Platforms

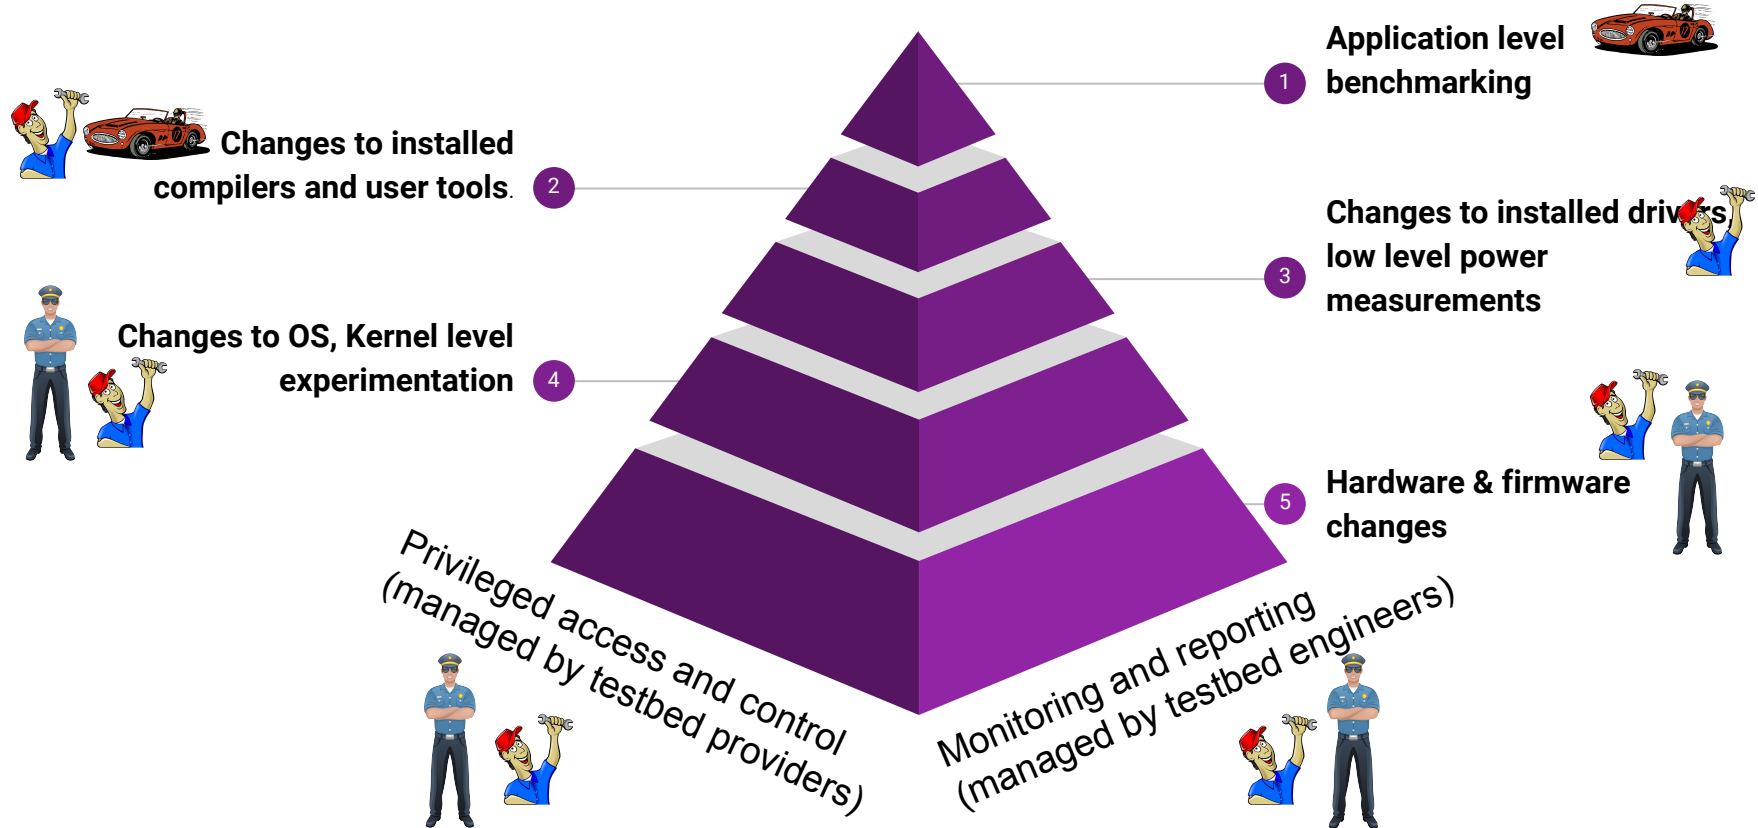
- ORNL NCCS Open Resources (2-factor authentication)
  - a. ARM1 system <https://www.olcf.ornl.gov/computing-resources/arm1/>
  - b. HPE-ARM system (system name unknown still)
  - c. Cray XC40 Knights Landing (Percival) 168 nodes
  - d. IBM Minsky (P8+/Pascal) (Summitdev) 54 node
- TokyoTech
  - a. TSUBAME KFC

Challenges: Technical & Policy

# Challenges

- Policy: Governance and management
  - Application process
  - Authentication and authorization
  - Access policies
- Technical: System provisioning and scheduling (isolation and security)
  - Innovative hardware ecosystem
  - Virtualization
  - Metal-as-a-service
  - Access to external environments e.g. quantum systems thru cloud

# Management of Privileged Access



# *Lightweight* Application Process for Accessing Testbeds

## Apply for Access

### Personal Information

#### Title

- None -

#### First Name / Given Name \*

#### Middle Name

#### Last Name / Surname \*

#### Street Address Title

If applicable, e.g. "Department of Computational Chemistry"

#### Street Address \*

#### Street Address (Cont.)

#### City

#### State/Province/Region \*

#### Zip/Postal Code \*

#### Country \*

### Project Information

#### Project PI

Enter the name of the project's principal investigator.

#### Detailed Description

Please include a detailed description of the work you will be doing. Include the names of major application codes you intend to use, along with descriptions of any programming models, software libraries and/or resources you will need. This should be 2-4 sentences long unless you are only applying for additional system access.

#### Existing Username

#### OCAMS/XCAMS Username \*

[Create an XCAMS username.](#)

#### Existing ExCL Username

☐ No

☐ Yes

Do you have an ExCL username from a previous/existing account?

#### Preferred Shell \*

### Comments

#### Comments or Questions?

Low overhead application process

# NEW USER OF EXISTING PROJECT

As soon as the project has been approved, Principal Investigators (PIs) should apply for their own account ("**New PI**" form).

Once they have an account they will be able to grant access to members of their research group. PIs will receive an e-mail request to confirm access for the group members. **Accounts will be opened only after receiving an e-mail of approval.**

## Low overhead application process

**Registration**

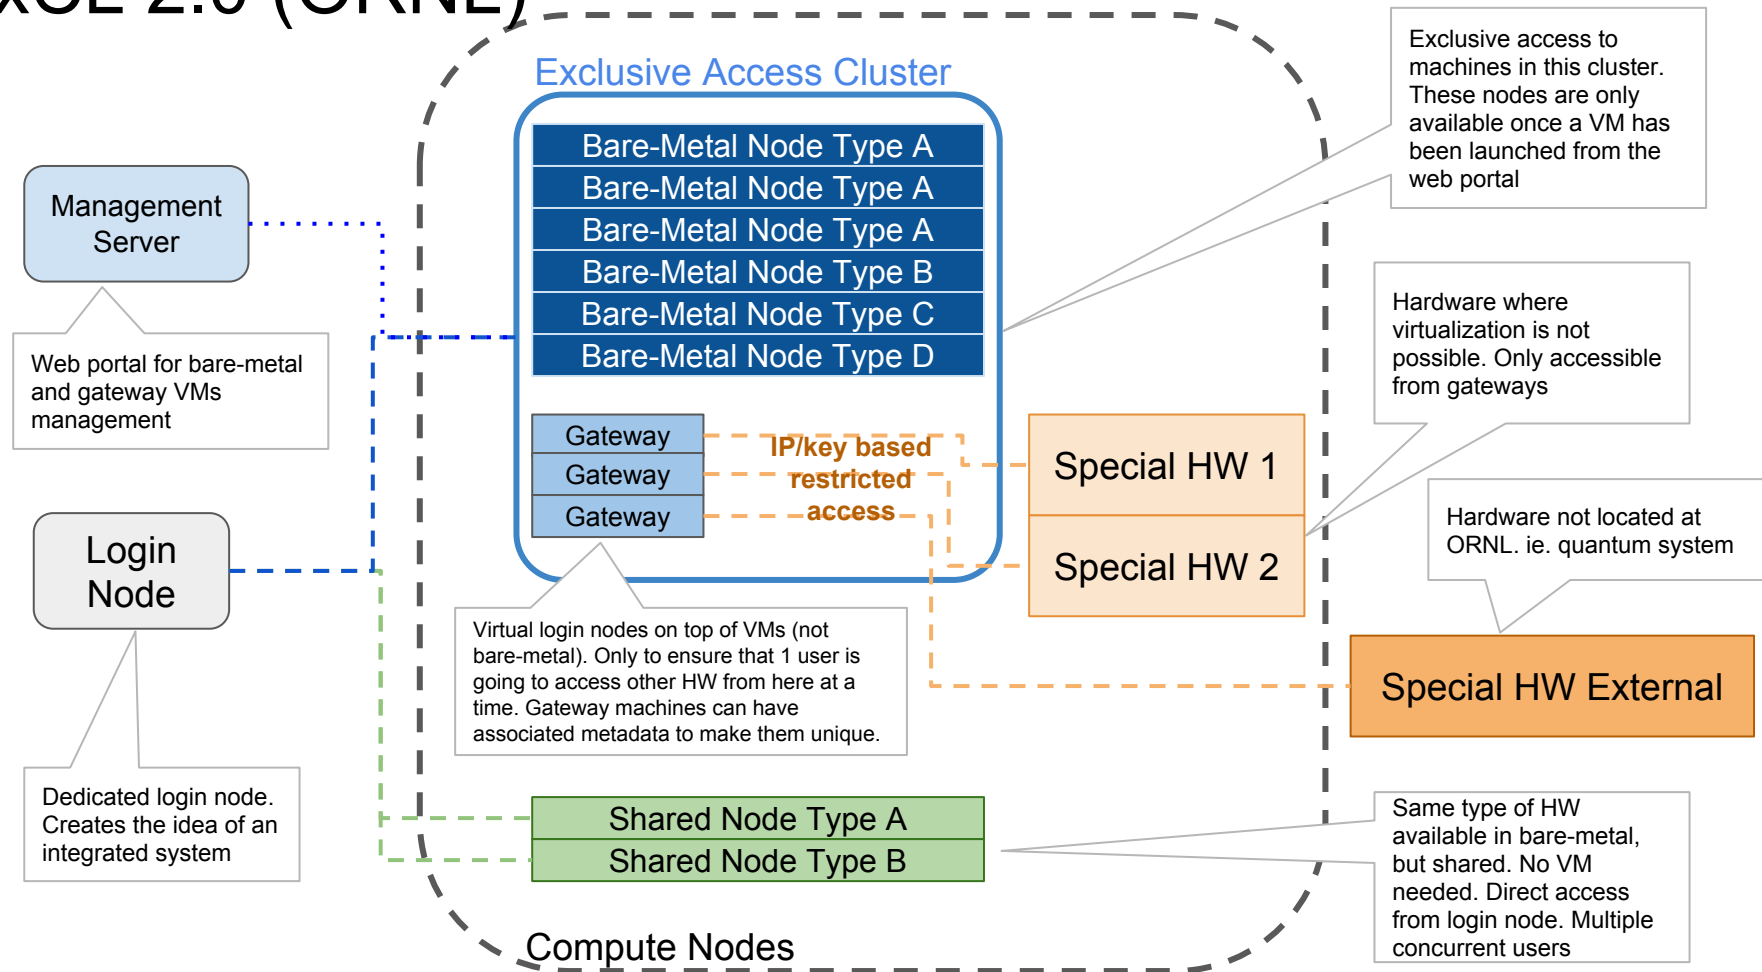
<b>Username proposal:*</b> <i>Username can be max 8 characters long, and must not contain blanks nor special characters</i>	<input type="text"/>
<b>Initial password:*</b> <i>First 4 letters. The rest of the password will be sent to the applicant after the account creation The resulting password must be changed at the first login</i>	<input type="password"/>
<b>Principal Investigator:*</b> <i>Last name First name &gt; Project ID</i>	<input type="text" value="Select PI"/>
<b>Title:*</b>	<input type="text" value="Mr."/>
<b>First name:*</b>	<input type="text"/>
<b>Last name:*</b>	<input type="text"/>
<b>Nationality:*</b>	<input type="text"/>

ORNL XCAMS & CSCS account application is similar

# Architecture

Addressing Experimental (incl. NDA) and  
Evaluation Platform needs

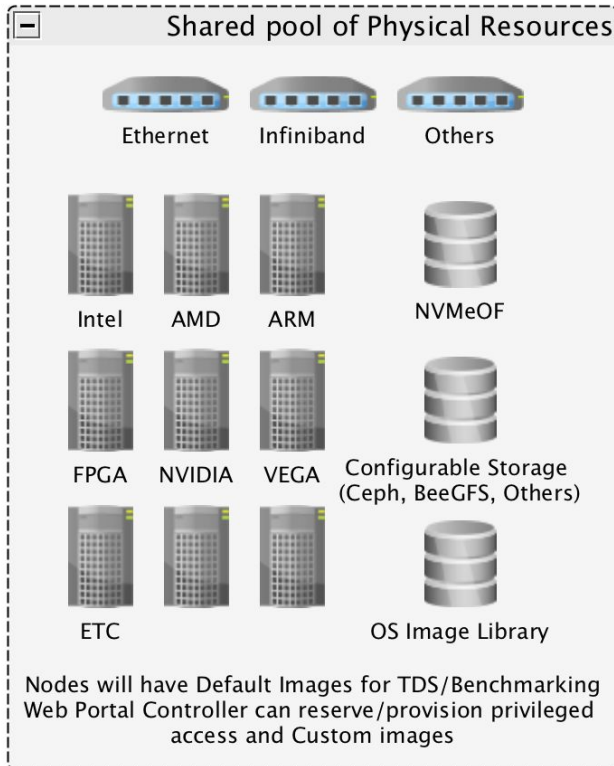
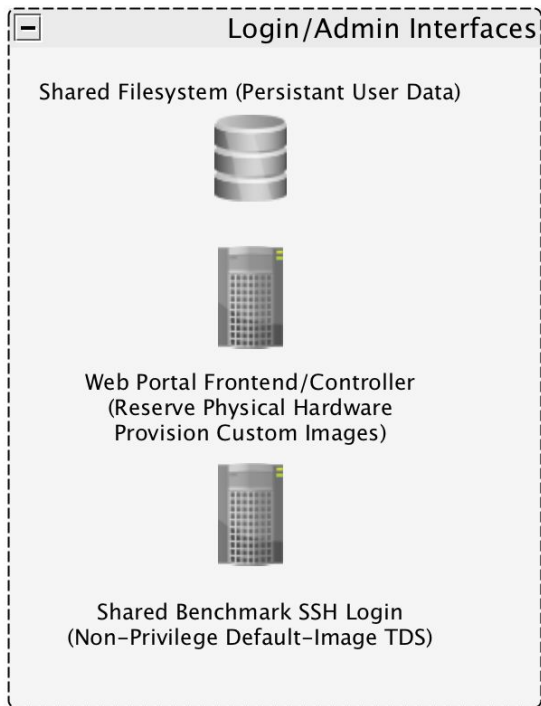
# ExCL 2.0 (ORNL)



# CSCS Testbed Architecture

Hardware connected to Monitored PDUs

(for power measurements)



Operation **default** for most nodes is **shared** access ([slurm w/container support](#)).

Allowed users can **re-provision** a server into a **dedicated** resource for various test cases: VMs, Containers, or Bare Metal deployments. [Clean up on release & return to shared mode.](#)

Users to be able to provide **own images** for **dedicated** modes, with some pre-made available for ease of use.

NDA hardware can be hidden from certain users/groups (not operated in Shared mode, only available for Dedicated Provisioning)

# Proposal for Federated AAI Ease of Access & Sharing of Resources

# Concept Similar to OpenID Connect

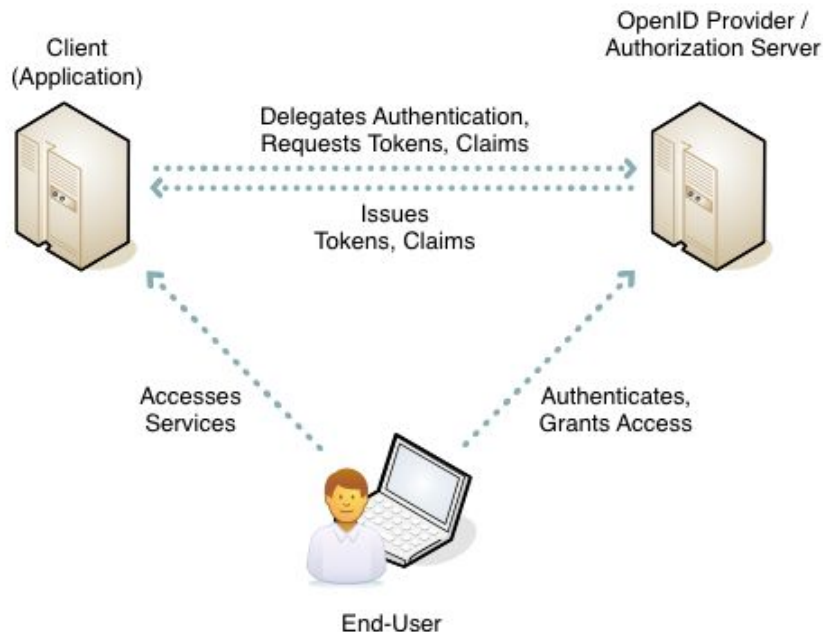
## Log in with OpenID

Click your OpenID account provider:



Or, manually enter your OpenID URL:

Or, if you don't have an OpenID through any of the above, [click here to sign up!](#)

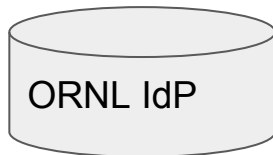


# ADAC Federation with Trusted Identities

- ADAC User and Project DB
  - ORNL ExCL and CSCS users initially
  - Proof of concept for web based technologies
- Identity brokering with KeyCloak



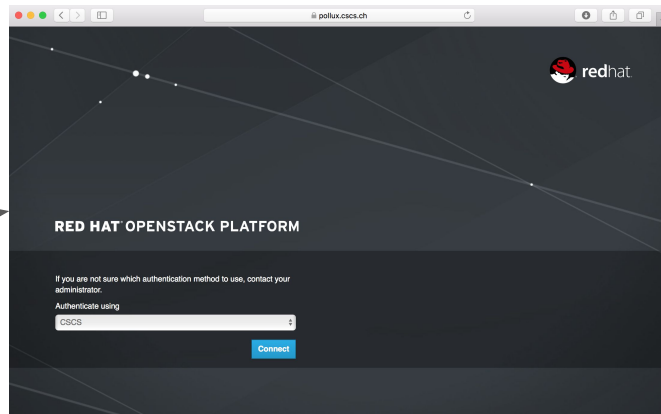
ORNL registered user



1. Access

4. Authorized to access  
CSCS service

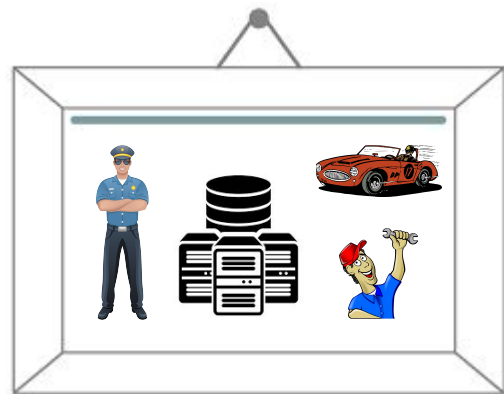
CSCS OpenStack Platform



2. Select ORNL IdP for  
authentication  
3. Redirected to ORNL

# Status & Next Steps

- Finalize federated AAI and central services decisions
  - Testing and implementation of Keycloak
  - ADAC user and proj DB central management
- Introduce a landing page for ADAC testbeds
  - Links to other sites
  - Inventory
- Tune testbed implementation for use cases
  - Enable HPC features
  - Metal-as-a-service
  - Verify functionality of OpenStack services
    - Bare Metal (Ironic)
    - VMs (Nova)
    - Container Orchestration (Magnum)
    - Resource Reservation (Blazar)



Vendors engagement is critical

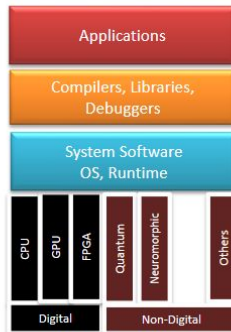
# DOE Workshop on Extreme Heterogeneity

23-25 Jan 2018

# Extreme Heterogeneity Workshop

Jan. 23-25, 2018, in Gaithersburg, MD

- **POC:** Lucy Nowell ([Lucy.Nowell@science.doe.gov](mailto:Lucy.Nowell@science.doe.gov))
- **Goal:** Define challenges that extreme heterogeneity presents to the software stack and programming environment and identify related Computer Science priority research directions that are essential to making extremely heterogeneous systems useful, usable and secure for science applications and DOE mission requirements in the 2025-2035 timeframe.
- 148 expected participants: DOE Labs, academia, & industry
- ~20 observers from DOE and other federal agencies (DoD, NSF, NASA)
- Pre-workshop report is being edited and will be posted by Jan. 1, 2018
- 105 white papers were received by the Dec. 4 deadline
  - After review, these resulted in 26 new invitations to Lab people and 20 to non-Lab people, including academics, industry and people from Europe and Japan.
- Agenda is being finalized, based in part on white paper content



## Tuesday, January 23, 2018

- 10:00 – 10:15 Introductions: Lucy Nowell and Jeffrey Vetter
  - 10:15 – 10:35 Welcome and ASCR Update – Barbara Helland, Director, Advanced Scientific Computing Research
  - 10:35– 11:05 View from ASCR Research Division – Steve Lee, Acting Division Director
  - 11:05 – 11:35 Invited Plenary Talk: IEEE Rebooting Computing - Tom Conte
  - 11:35– 11:45 Break
  - 11:45 – 12:15 Invited Talk: Architectural Trends and System Design Issues - Bob Colwell
  - 12:15– 1:30 FSD Introduction to Extreme Heterogeneity – Jeffrey Vetter, John Shalf, and Maya Gokhale + FSD section owners
  - 1:45 – 2:30 Break for lunch
  - 2:30 – 3:15 Invited Talk: Report on the ASCAC future of computing study - Maya Gokhale
  - 3:15 – 4:30 Panel on Issues Raised by Extreme Heterogeneity - Moderator Ron Brightwell
- Usability, Understandability and Programmability – Salman Habib
- Operating and Runtime Systems – Ron Brightwell
- Data Analytics – Wes Bethel
- EH Workflow Management – Ewa Deelman
- Open Q&A

# Status

- Gov shutdown forced cancellation of physical meeting
- Moved to virtual meeting
  - Kept to original agenda (with some minor changes for timezones)
  - Approximately 200 participants viewed plenary sessions!!
- Breakout groups converged on priority research directions

BOG ID	B/O Topic	Tu PM	Wed PM-1	Wed PM-2	Th AM	Th PM
1	Prog Env: Abstractions, Models, and Languages	Aiken/McCormick			McCormick	
2	Data Management and I/O	Ross/Byna				Ross
3	Data Analytics and Workflows	Tom P./Yoo		Christine S./Bethel		
4	OS/RM: global, composition, workflow		Brightwell			
5	Software Development Methodologies		Li/Bernholdt			
6	Crosscut: Modeling and Simulation		Chien/Donofrio/Leidel		Wilke/Lan/Gokhale	
7	Prog Env: Compilers, Libraries, and Runtimes			Strout/Chapman		
8	System Management, Admin, Job Scheduling			Peltz/Hartman-Baker		
9	Crosscut: productivity, composition, interoperability			Lucas		
10	OS/RM: local, prog env support				Lang	
11	Crosscut: Portability, code reuse, performance portability				Dubey/Li	
12	Prog Env: Debugging and Correctness, autotuning, specialization					Hall/Mellor-Crummey
13	Crosscut: resilience, power					Cappello/Cameron

# Status

- Initial Priority Research Directions (Categories)
  - Programmability and Software Development Productivity
  - Managing Execution, Scheduling
  - Correctness, Debugging, Reproducibility
  - Modeling and Simulation for Performance, Power, Resiliency
- Deliverables
  - Slides w/ PRDs be the end of the workshop / Done
  - Brief report for HQ by Mar 1
  - Final public report by May 1