

# Yellowstone early experience and access schedule

David Hart, CISL/USS

Oct. 18, 2012

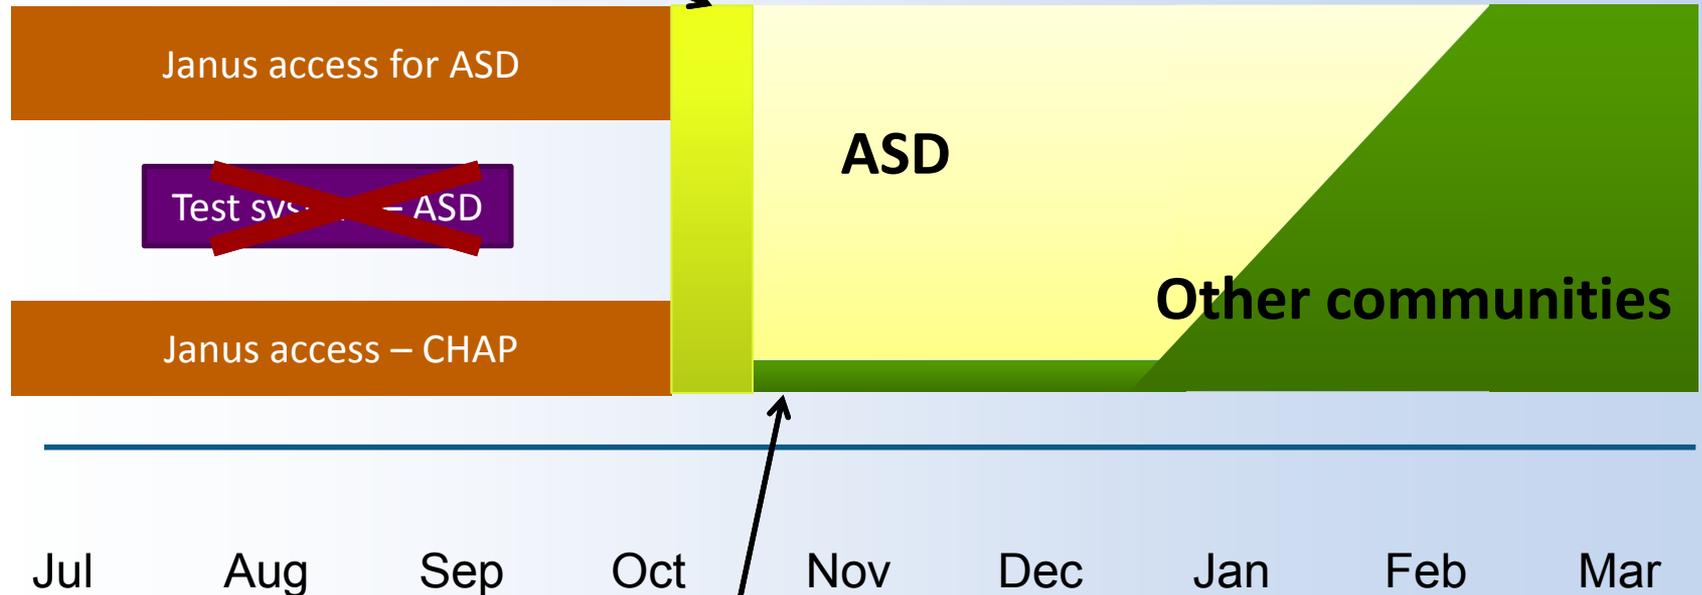
# Yellowstone Schedule

- **Initial production target: June 2012**
  - Delayed due to decision to go with full FDR
- **Hardware all delivered by end of June 2012**
  - First powered up mid-July
- **Initially expected beginning of September 2012 for production**
  - InfiniBand performance challenges delayed early user access until October 8
- **About 50 ASD users given early access on October 8**
  - Along with ~50 other NCAR early users, including CESM and WRF developers
- **Targeting October 24 for official ASD start**
  - Following final reconfiguration of GLADE management environment
- **November 1 for access by general user population**

# Yellowstone access schedule

Early use: Oct. 8–24

## Yellowstone



4% of Yellowstone (~1 Bluefire-equivalent) available for other communities to begin porting and testing work during ASD.

# Yellowstone

## Accelerated Scientific Discovery (ASD)

- **Projects represent 16 codes**
  - **CESM's CAM5, POP2, CLM4 & CICE; MPAS; WRF-Chem; MITgcm; ECMWF Ensemble Prediction System**
  - **NCAR LES; GHOST; Hybrid DNS & Bin EULAG; HIPSTR**
  - **AWP-ODC & SPECFEM3D; P3D**
- **Other codes/users already running**
  - **MURaM** (Matthias Rempel, HAO), **MFE** (Yuhong Fan, HAO), **FWSDA** (Po Chen, UW)
- **Scaling targets for ASD projects (max planned job sizes)**
  - 65,536 cores or more: 1
  - 40,000-44,200 cores: 2
  - 16,384 cores: 4
  - 4,096 cores: 1
  - 2,048 cores: 3

# University ASD projects

- **James Kinter, COLA — 21 million core-hours**
  - Towards seamless high-resolution prediction at intraseasonal and longer timescales (ECMWF models)
- **Lance Collins, Cornell U — 19 million core-hours**
  - DNS of cumulus cloud core processes over larger volumes and for longer times (HIPSTR)
- **B. Fox-Kemper, CU-Boulder — 16 million core-hours**
  - Arrest of frontogenesis in oceanic submesoscale turbulence (MITgcm)
- **Thomas Jordan, USC — 7.3 million core-hours**
  - Community computational platforms for developing 3-D models of earth structure (AWP-ODC, SPECFEM3D)
- **Michael Shay, U Delaware — 7.2 million core-hours**
  - Turbulence in the heliosphere: The role of current sheets and magnetic reconnection (P3D)

# NCAR ASD projects

- **Gabriele Pfister, ACD — 6.25 million core-hours**
  - Prediction of North American air quality (WRF-Chem, CAM5)
- **Annick Pouquet, IMAGE — 11 million core-hours**
  - Rotation and stratification at unit Burger number in turbulent flows (GHOST)
- **David Richter, MMM — 5.2 million core-hours**
  - Turbulence modification in the spray-laden atmospheric marine boundary layer (NCAR LES)
- **William Skamarock, MMM — 6.5 million core-hours**
  - Global cloud-permitting atmospheric simulations using MPAS (MPAS)
- **Justin Small, CGD — 25.2 million core-hours**
  - Meso- to planetary-scale processes in a global ultra-high resolution climate model (CESM—CAM5, POP2, CLM, CICE)
- **Andy Wyszogrodzki, RAL — 6.25 million core-hours**
  - Petascale simulation of physics and dynamics of turbulent clouds (HDNS, Bin EULAG)

# Training for Yellowstone

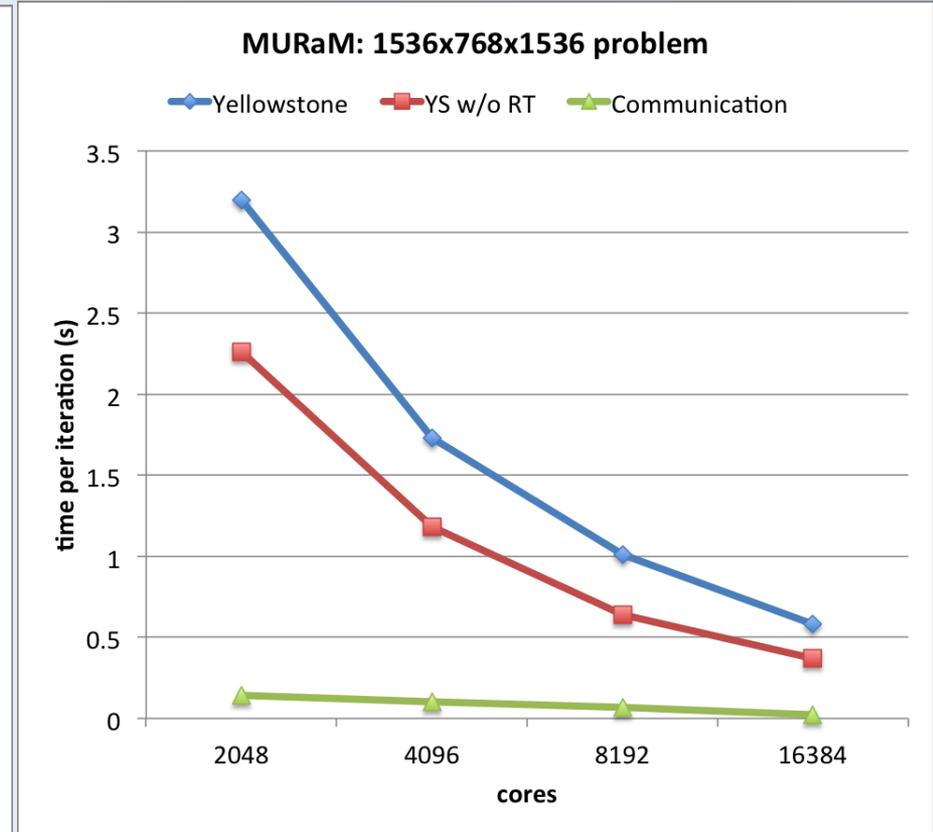
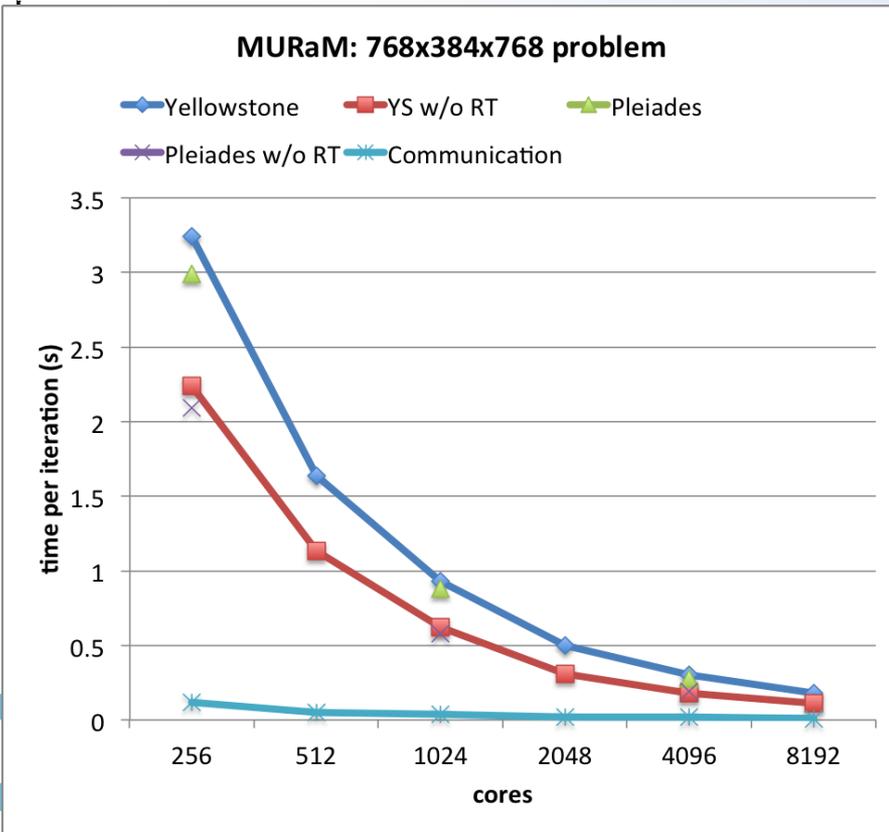
- **“Getting Started with Yellowstone”**
  - On-site for NCAR users
  - Via Adobe Connect for University users
  - Held Oct. 10 for ASD/early users
  - Additional dates to be set in Fall and Winter
- **Transition documentation**
  - <https://www2.cisl.ucar.edu/resources/yellowstone/transition>
  - Much documentation has been or will be updated for Yellowstone
    - <https://www2.cisl.ucar.edu/docs/newusers>

# Early user experiences

- **Users with portable and scalable codes were quickly able to conduct scaling tests up to 16,384 cores**
  - Beyond 16k cores, IB performance noticeably affected job performance and success rates
  - Compilation environment working well
- **“By the way, Yellowstone is performing excellent with everything I've thrown at it. Even a bit faster than I had expected.”**
  - David Richter, MMM
- **“My code has run with up to 3072 cores without problems. It seems to be about 25% faster than what I get in running it on the same Sandy Bridge nodes on Pleiades.”**
  - Yuhong Fan, HAO

# Example: MURaM scaling

atory



Com

# However...not done yet

- **InfiniBand performance remains less than optimal, particularly at scale**
  - CISL, IBM, Mellanox held a “Routing Summit” yesterday to address the issue
- **Major changes on Friday, to IB routing and the latest update of IBM Parallel Environment (PE)**
  - System performance and job success in the process of being resolved
- **Access to HPSS next week**
  - HSI, HTAR to be installed

# Going forward

- **Work on GLADE going on behind the scenes**
  - In advance of “multi-cluster” configuration change next week
- **Charles Archer, IBM application expert, is on site at Mesa Lab through next week to work with CSG and users to improve system stability and usability for applications**
- **IBM is working with CISL to ensure that all ASD applications can run as planned**
- **Targeting Oct 24 to officially put ASD projects on the clock.**

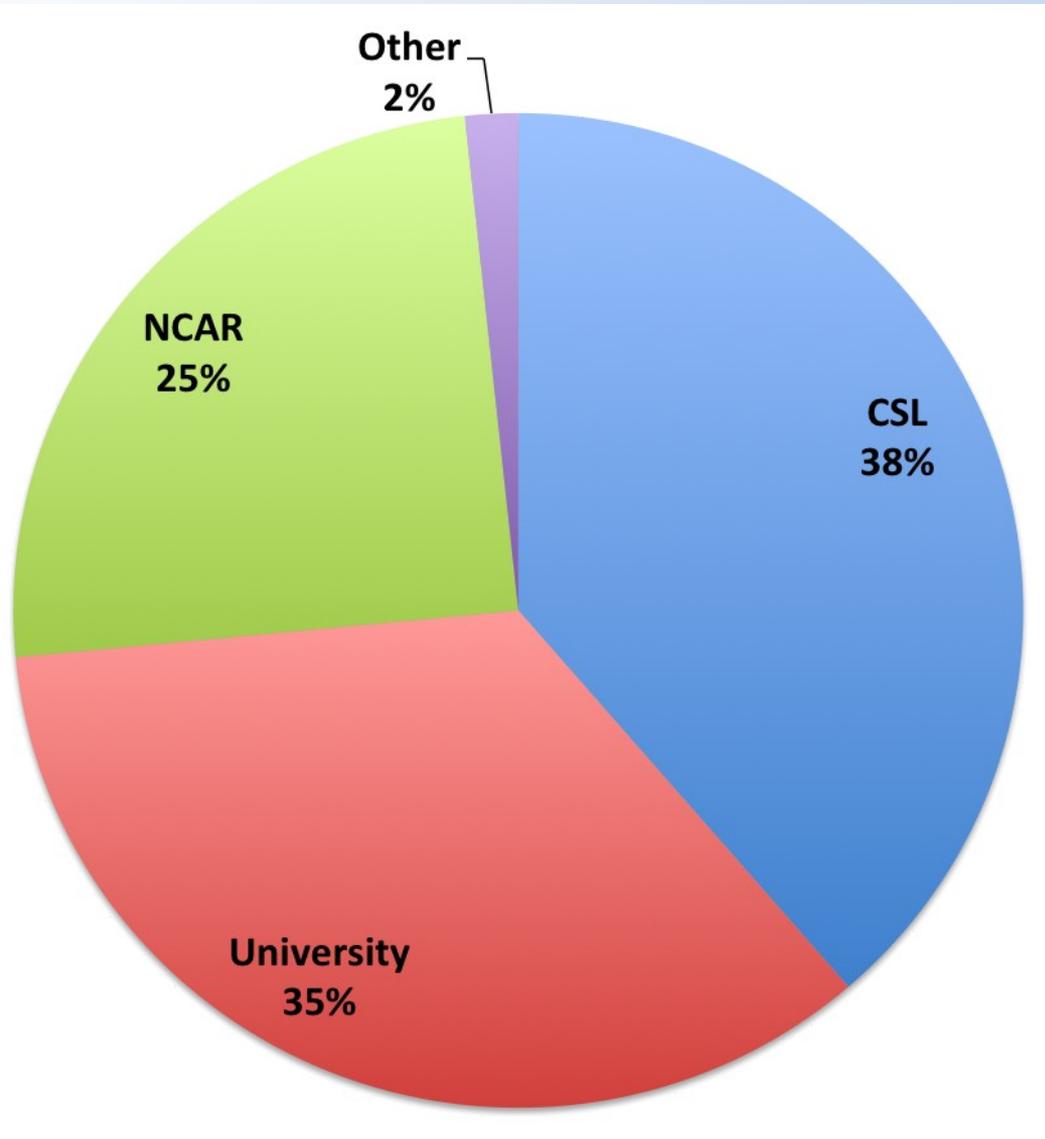
# Bluefire update

- **Extension awards made**
  - CSL projects, NCAR labs extended
  - University projects granted up to 200k GAUs with justification
  - University projects granted additional 30k GAUs upon request (normal policy)
- **No new projects initiated after Sept. 1, 2012**
- **Anticipate Bluefire available until at least Dec. 31, 2012**
- **Plan: As of Nov. 1 (or when Yellowstone available to all users), we will let users run without usage throttles**
  - Use will still be balanced among CSL, NCAR, University

# University projects requesting extra Bluefire GAUs

- Allison Wing
  - Steve Vavrus
  - Elie Bou-Zeid
  - David Reusch
  - Louisa Molina
  - Anna Barros
  - James Randerson
  - Lian-Ping Wang
  - Eric Chassignet
  - Chien Wang
  - Yang Zhang
  - Young-Oh Kwon
  - Charles Meneveau
  - Joyce Penner
  - Seth Olsen
- **Estimated GAUs available during July-October: 3.6 million**
  - **Awards made (greater than 30k): 2.2 million**

# Bluefire use, FY12





System Accounting Manager

**SAM**

# Managing the transition

- **Many (but not all) projects are being migrated from Bluefire to Yellowstone**
  - We are using the transition to eliminate inactive projects and user accounts, as well as unnecessary legacy system and user info
- **Bluefire to Yellowstone accompanied by accounting system transition (ACC8 → SAM)**
  - New capabilities and interfaces
  - Considerable effort from USS, WEG & HSS to automate processes, streamline tasks, and provide a consistent user environment across resources

# New system, new accounting

- **Current accounting system (“ACC8”) has been showing its age and proving difficult to maintain**
- **SAM is a modernized system with many new features and user-oriented reporting possibilities**
  - Web-based interface for data entry and reporting
    - Replaces Oracle Forms-based data entry and CISL Portal-based reporting
  - First production release will have basic functionality, with plans to expand and add new features

# SAM enhancements

- **Manages user accounts and access via LDAP**
- **Separate resource allocations**
  - To help monitor and track use of data analysis and visualization clusters, storage resources
  - Encompasses Janus allocations and usage
- **New alphanumeric project codes**
  - Used across systems for job charging and other purposes (project spaces, HPSS, etc.)
- **User-controlled preferences**
- **Multi-level reporting for project usage**

**Live demo (test system)**

**SAM Demo**



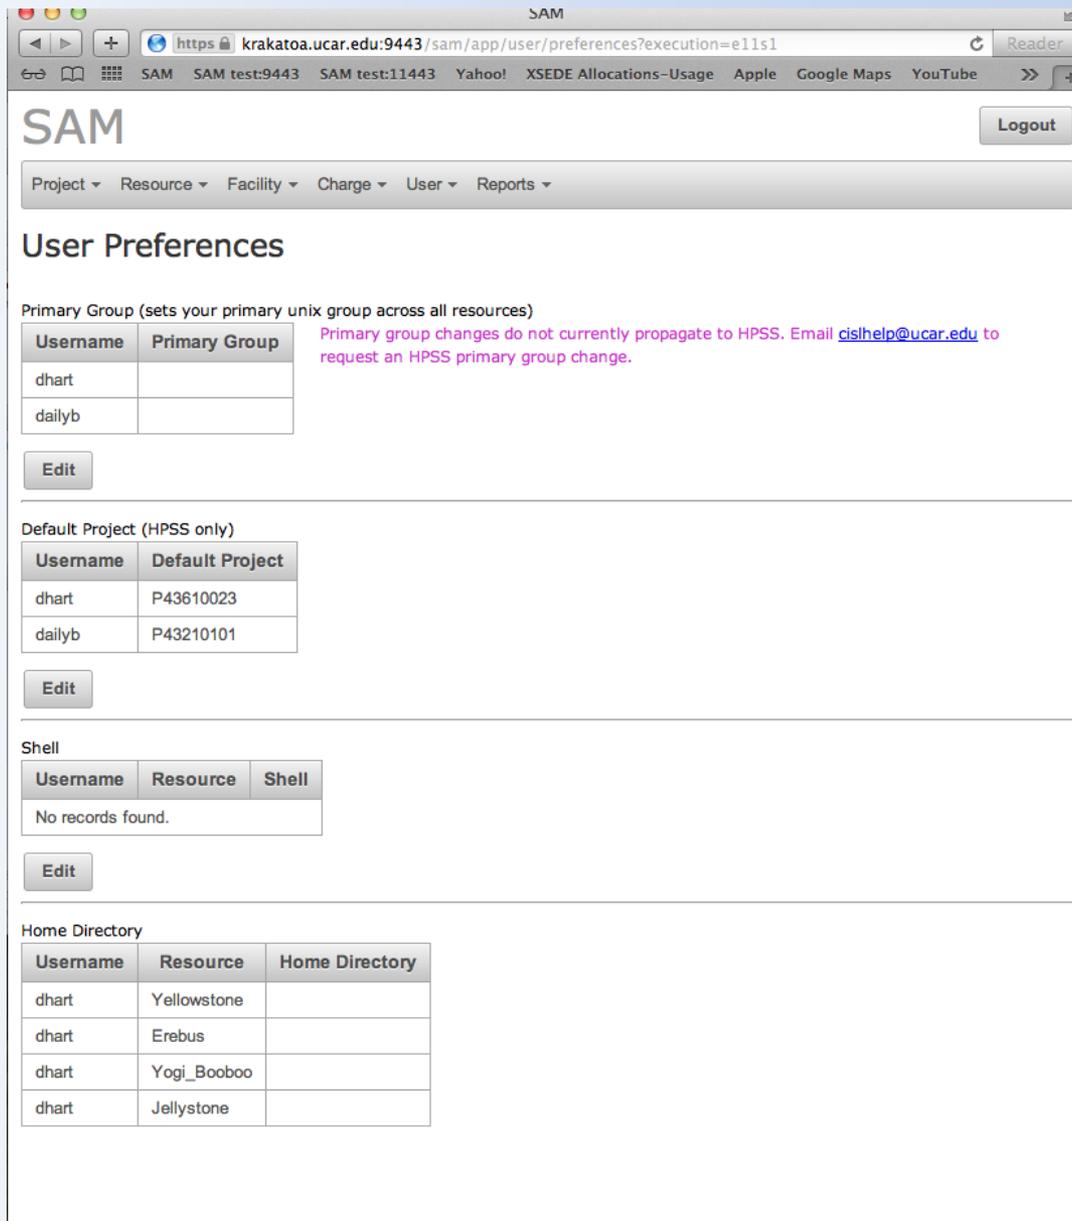
# QUESTIONS?

## SAM User Preferences

A user can view and modify certain system settings from within SAM.

This screen replaces a ticket request (primary group change), a visit to the CISL Portal (HPSS default project change), and a command-line utility (shell change).

Users can view, but not change their home directories.



The screenshot shows the SAM web interface for user preferences. The browser address bar shows the URL: `https://krakatoa.ucar.edu:9443/sam/app/user/preferences?execution=e11s1`. The page title is "SAM" and there is a "Logout" button in the top right corner. Below the title is a navigation bar with dropdown menus for "Project", "Resource", "Facility", "Charge", "User", and "Reports".

### User Preferences

Primary Group (sets your primary unix group across all resources)

Primary group changes do not currently propagate to HPSS. Email [cislhelp@ucar.edu](mailto:cislhelp@ucar.edu) to request an HPSS primary group change.

Username	Primary Group
dhart	
dailyb	

---

Default Project (HPSS only)

Username	Default Project
dhart	P43610023
dailyb	P43210101

---

Shell

Username	Resource	Shell
No records found.		

---

Home Directory

Username	Resource	Home Directory
dhart	Yellowstone	
dhart	Erebus	
dhart	Yogi_Boofoo	
dhart	Jellystone	

## SAM Account Statement (1)

Users will go to SAM for reports on project activity.

Upon logging in, they will see their list of projects.

To see project activity, select a project and click "Detail."

The screenshot shows a web browser window with the URL <https://tkrakatoa.ucar.edu:11443/sam/app/report/accountStatement/user?execution=e10s1>. The page title is "SAM" and there is a "Logout" button in the top right corner. Below the title is a navigation bar with dropdown menus for "Project", "Resource", "Facility", "Charge", "User", and "Reports". The main heading is "Account Statement". Below this is a table titled "Projects" with two columns: "Project Code" and "Title". The table contains three rows of data. Below the table is a "Detail" button.

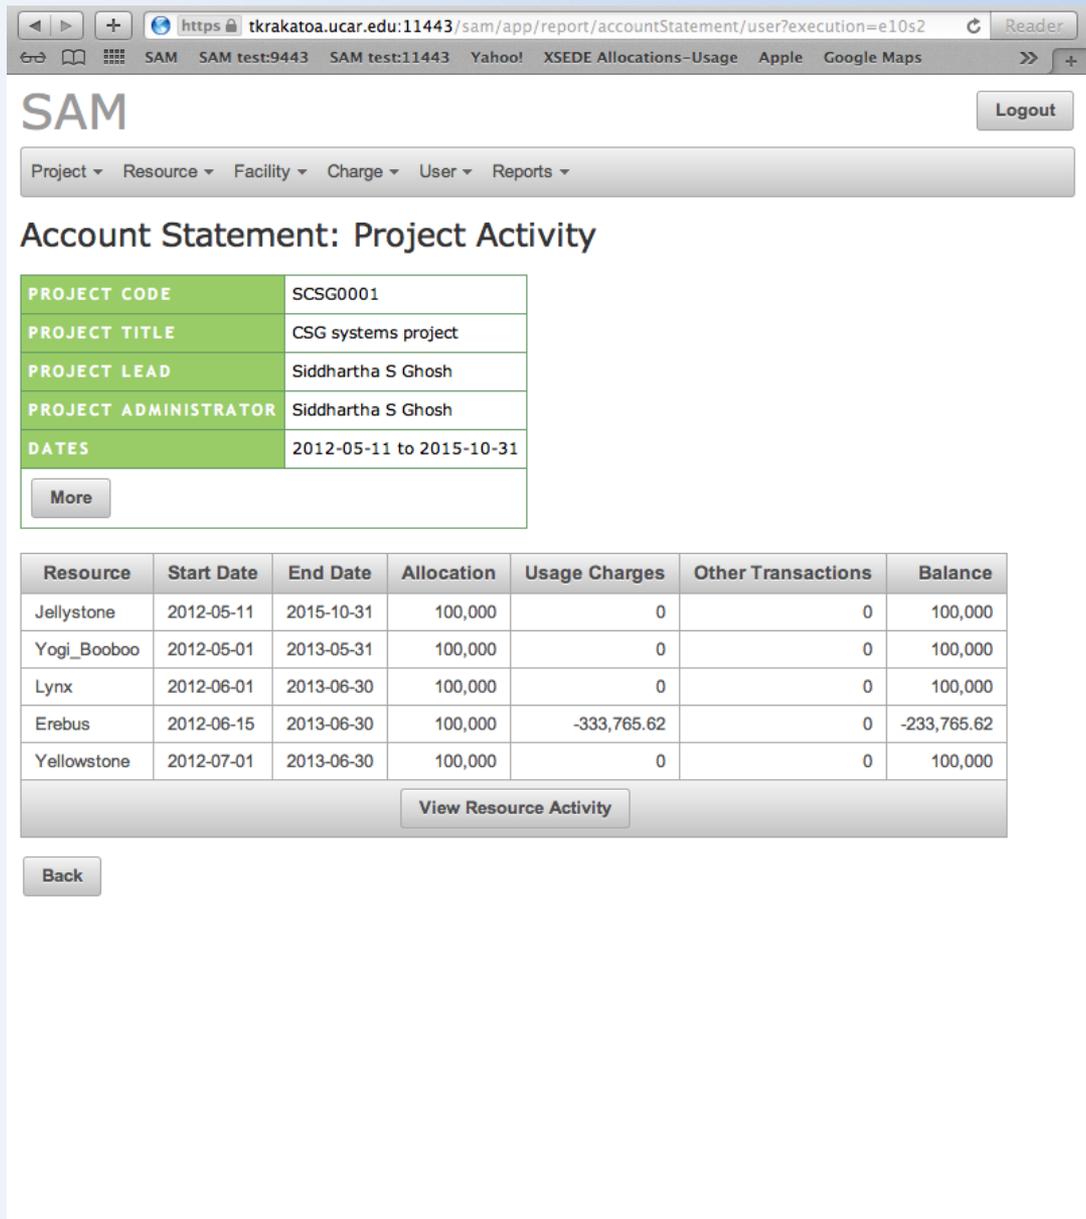
Projects	
Project Code	Title
SCSG0001	CSG systems project
P43610023	Consulting Services Group
UCSG0001	testing Data_Access and acct.txt

## SAM Project Activity (2)

The project activity report shows the projects current allocations (and link to prior allocations, if applicable) along with current charges, other transactions, and the current balance on those allocations.

User can select an allocation and “View resource activity.”

Two levels of detail are then available.



The screenshot shows a web browser window with the URL `https://tkrakatoa.ucar.edu:11443/sam/app/report/accountStatement/user?execution=e10s2`. The page title is "SAM" and there is a "Logout" button in the top right corner. Below the title is a navigation bar with dropdown menus for "Project", "Resource", "Facility", "Charge", "User", and "Reports". The main heading is "Account Statement: Project Activity".

Below the heading is a table with the following data:

PROJECT CODE	SCSG0001
PROJECT TITLE	CSG systems project
PROJECT LEAD	Siddhartha S Ghosh
PROJECT ADMINISTRATOR	Siddhartha S Ghosh
DATES	2012-05-11 to 2015-10-31

There is a "More" button below this table.

Below the "More" button is a table with the following data:

Resource	Start Date	End Date	Allocation	Usage Charges	Other Transactions	Balance
Jellystone	2012-05-11	2015-10-31	100,000	0	0	100,000
Yogi_Booboo	2012-05-01	2013-05-31	100,000	0	0	100,000
Lynx	2012-06-01	2013-06-30	100,000	0	0	100,000
Erebus	2012-06-15	2013-06-30	100,000	-333,765.62	0	-233,765.62
Yellowstone	2012-07-01	2013-06-30	100,000	0	0	100,000

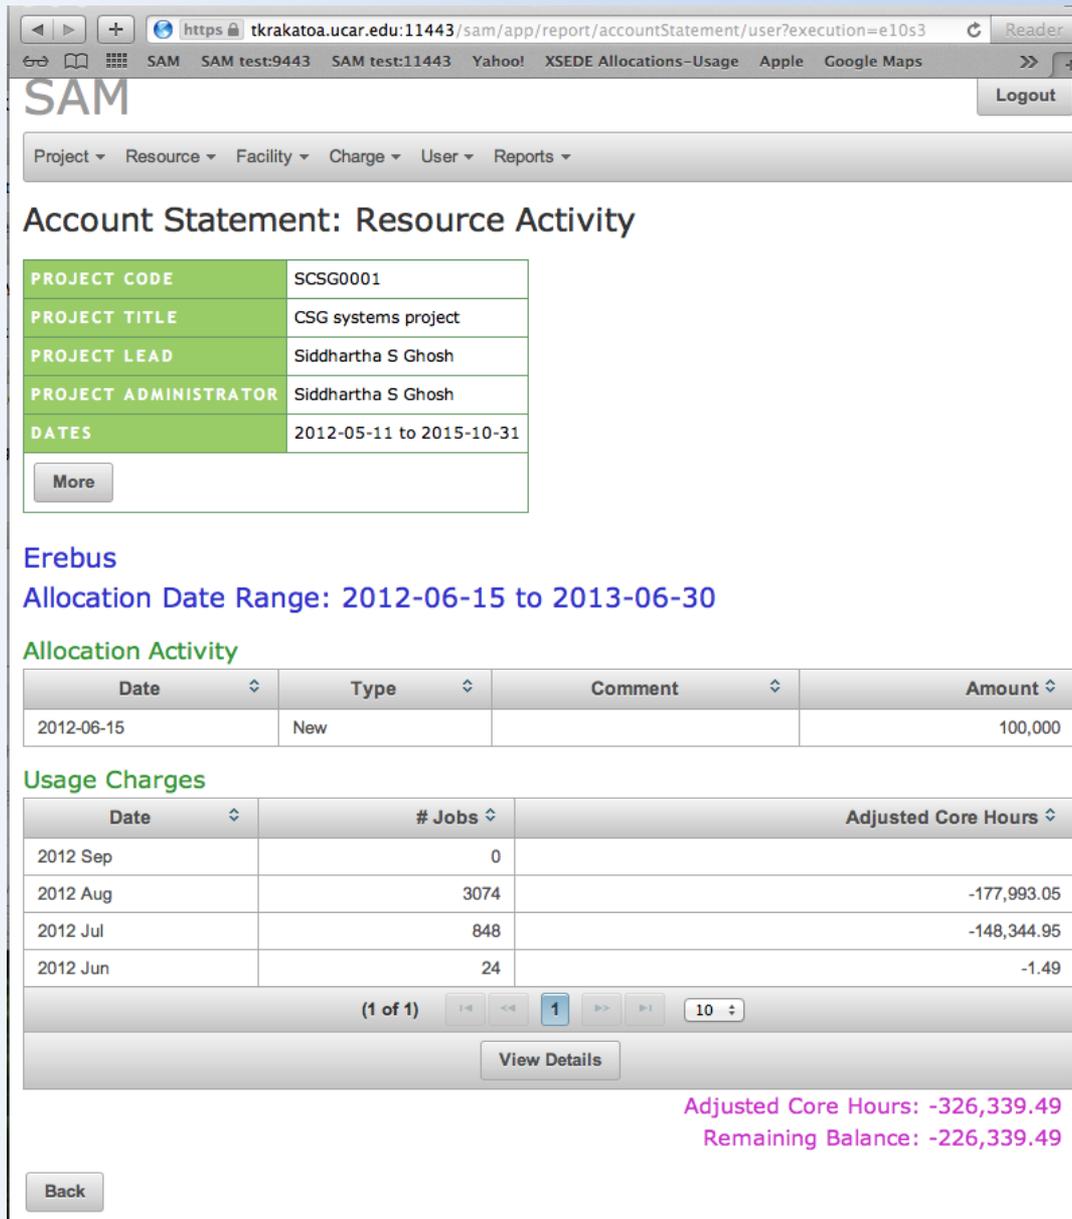
Below this table is a "View Resource Activity" button. At the bottom left of the page is a "Back" button.

## SAM Resource Activity (3)

For compute allocations, SAM will show allocation activity, including the original award as well as any subsequent modifications (debits, supplements, etc.)

For compute charges, SAM shows monthly roll-up of charges and the number of jobs, across all users on the project.

For all the details, user can select a month and will be provided job-level details for the month, for viewing or download to .csv file.



The screenshot shows the SAM web application interface. At the top, there is a navigation bar with the SAM logo and a 'Logout' button. Below the navigation bar, there are dropdown menus for 'Project', 'Resource', 'Facility', 'Charge', 'User', and 'Reports'. The main content area is titled 'Account Statement: Resource Activity' and displays a table with the following data:

PROJECT CODE	SCSG0001
PROJECT TITLE	CSG systems project
PROJECT LEAD	Siddhartha S Ghosh
PROJECT ADMINISTRATOR	Siddhartha S Ghosh
DATES	2012-05-11 to 2015-10-31

Below the table is a 'More' button. Underneath, there is a section for 'Erebus' with an 'Allocation Date Range: 2012-06-15 to 2013-06-30'. This is followed by an 'Allocation Activity' table:

Date	Type	Comment	Amount
2012-06-15	New		100,000

Next is a 'Usage Charges' table:

Date	# Jobs	Adjusted Core Hours
2012 Sep	0	
2012 Aug	3074	-177,993.05
2012 Jul	848	-148,344.95
2012 Jun	24	-1.49

At the bottom of the usage charges table, there is a pagination control showing '(1 of 1)' and a 'View Details' button. To the right of the table, there are two lines of text in pink: 'Adjusted Core Hours: -326,339.49' and 'Remaining Balance: -226,339.49'. A 'Back' button is located at the bottom left of the page.

## SAM Resource Activity (4)

For storage allocations, the resource activity report shows a summary per storage accounting run – weekly for HPSS.

The summary shows the number of files, storage holdings in terabytes, as well as the allocation charges.

The detail report shows the storage per user on the project.



The screenshot shows the SAM web interface. At the top, there's a browser address bar with the URL `https://tkrakatoa.ucar.edu:11443/sam/app/report/admin?execution=e11s8`. Below the browser, the SAM logo is visible, along with a "Logout" button. A navigation bar contains dropdown menus for "Project", "Resource", "Facility", "Charge", "User", and "Reports".

The main heading is "Account Statement: Resource Activity". Below this is a table with the following data:

PROJECT CODE	P93300019
PROJECT TITLE	Cesmp-Hpss Only
PROJECT LEAD	Marika Holland
PROJECT ADMINISTRATOR	
DATES	2011-10-04 to 2012-11-01

There is a "More" button below the table. Below the table, the text "HPSS" is displayed, followed by "Allocation Date Range: 2011-07-01 to 2012-11-01".

The "Allocation Activity" section contains a table with the following data:

Date	Type	Comment	Amount
2012-07-24	New		100,000

The "Usage Charges" section contains a table with the following data:

Date	Files	Holdings (TB)	TB Years
2012-03-25	90,382	12.18	-0.25
2012-03-18	86,285	11.10	-0.23
2012-03-11	108,829	19.25	-0.40
2012-03-04	150,923	14.69	-0.31

Below the Usage Charges table, there is a pagination control showing "(1 of 1)" and a "View Details" button. At the bottom right, there is a summary: "TB Years: -1.19" and "Remaining Balance: 99,998.81". A "Back" button is located at the bottom left of the interface.