

# Ohio Supercomputer Center

An **OH·TECH** Consortium Member

OSC OnDemand: A Web Platform Integrating Access to  
HPC Systems, Web and VNC Applications

Dave Hudak, Thomas Bitterman, Patricia Carey, Douglas Johnson,  
Eric Franz, Shaun Brady, Piyush Diwan



# Overview

- Introduction and OnDemand Goals
- OnDemand User Experience
- OnDemand Implementation
- OnDemand Usage
- Conclusions and Future Work



# Introduction

- Web has become the dominant (I would say sole) mechanism for remote compute access in every area except HPC
  - Text content (blogs, articles, magazines, newspapers)
  - Audio (Pandora, Spotify, etc.)
  - Video (YouTube)
  - Banking, travel, shopping, course registration, ...
- And it's all available through the browser
  - Decreased use of Java applets, JNLP and flash
  - Rising interest in HTML5, PHP and Javascript
  - Support smartphones, tablets and (even) computers



# Existing Web Applications in HPC

- Lots of innovative, interesting applications
  - Science Gateways - designed to raise level of abstraction for user by incorporating a workflow
  - XSEDE User Portal – integrating file access, terminal access, documentation and more
  - HubZero and its products (notably nanoHub) – giving users the ability to run apps, write apps and publish apps
- Our group has built science gateways hosted in community accounts with application-level authentication and authorization
  - “HPC User” vs. “Web user”
- Wanted to build a general-purpose solution for HPC access that could support system access, web and viz apps



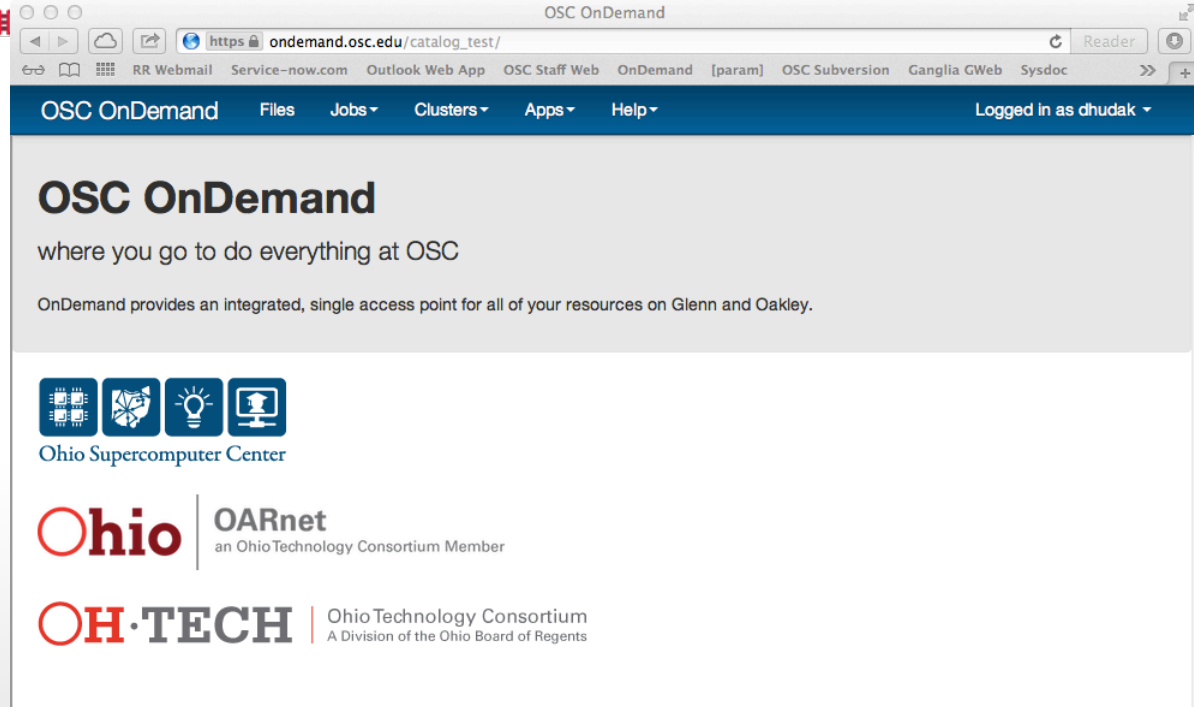
# OSC OnDemand Goals

- Single point of entry for OSC services
  - File access
  - Job management
  - Terminal access
  - Visualization apps
  - Web apps
- User needs three things
  - [ondemand.osc.edu](http://ondemand.osc.edu)
  - Username & password
- Completely browser-based
  - Support computers, tablets and smartphones
- Single sign-on
- Zero Install
  - No client-side software needed, apart from browser
- Firewall friendly
  - Keep traffic on https



# OnDemand User Experience

- User logs in with HPC credentials
- Presented with dashboard
- Authentication via OpenID, backed by LDAP

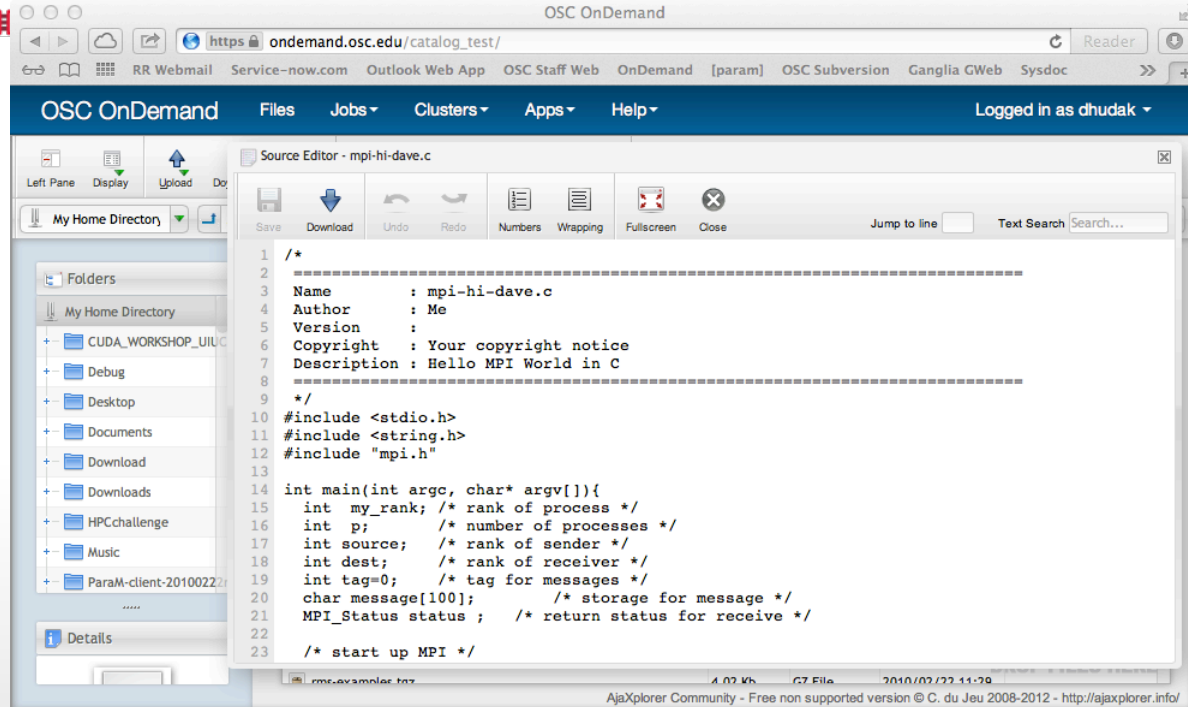


- Dashboard is top-level view to services
- Updated dashboard has web feedback functionality



# File Access

- AjaxExplorer is an open-source HTML5 app
- File system browser
- Upload and download files
- View image files
- Edit text files



- App is provided by per-user web server (PUWS)
- Apache instance launched with user, not system, privileges
- Any files created are owned by the user who logged in



# Job Constructor

- Generalization: “99% of all jobs are copies of other jobs”
- Template-based approach
  - Runnable template (think “Hello World”)
  - Best practices (I/O staging)

Date	Name	Cluster	Script	Job status	PBS ID	run
03/04/2013 - 10:05	OMP test	oakley	/ifs/07/dhudak/crimson_files/JobConstructor/58/script.sh	Not Submitted		58
02/28/2013 - 15:08	Patricia test	oakley	/ifs/07/dhudak/crimson_files/JobConstructor/57/script.sh	Completed	722358.oak-batch.osc.edu	57
02/28/2013 - 13:58	P2 test	glenn	/ifs/07/dhudak/crimson_files/JobConstructor/56/script.sh	Completed	9424616.opt-batch.osc.edu	56
02/28/2013 - 13:55	Patricia test	oakley	/ifs/07/dhudak/crimson_files/JobConstructor/55/script.sh	Not Submitted		55
02/28/2013 - 09:24	New test	oakley	/ifs/07/dhudak/crimson_files/JobConstructor/54/script.sh	Not Submitted		54
01/23/2013 - 15:46	Test 1	oakley	/ifs/07/dhudak/crimson_files/JobConstructor/53/script.sh	Not Submitted		53
12/21/2012 - 13:22	Tom test 3	oakley	/ifs/07/dhudak/crimson_files/JobConstructor/52/script.sh	Not Submitted		52

- User instantiates job template
- Newly created job has its own directory in the filesystem
- Links provided to open job in File view
- Once job customized, click “Submit”






# Job Monitor

- Active job
  - Queued or running
- Job view
  - Active jobs on either cluster
  - Sortable via username, job name, job status, date, cluster
  - Search by username

opt2159 opt2180 opt2238 opt2404

## CPU Report

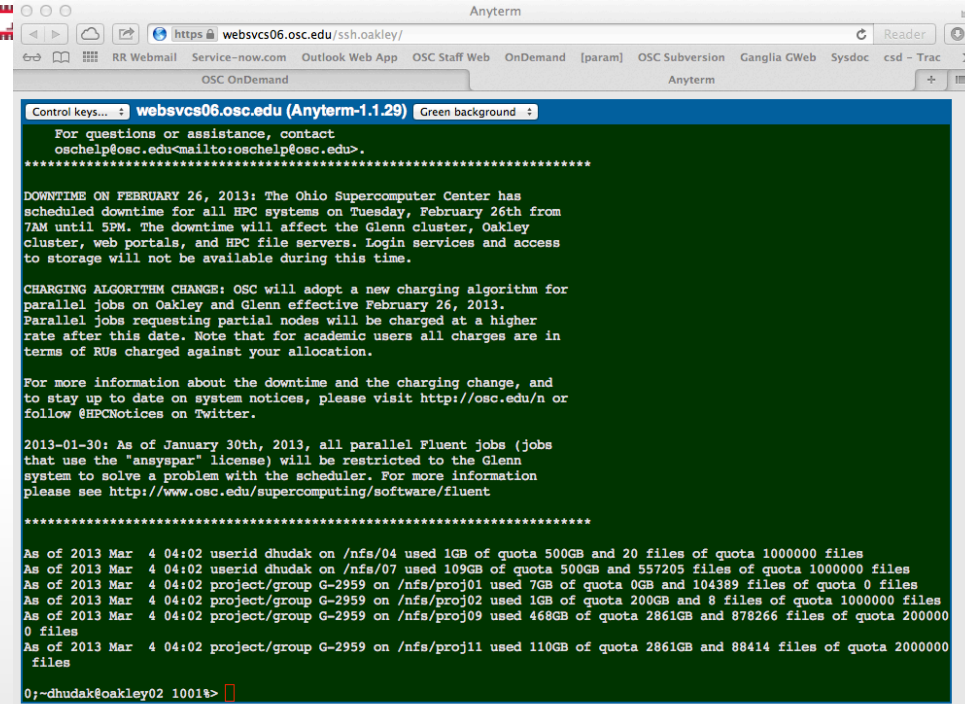


- Individual node statistics on a per-job basis
  - Above, CPU load for 1 of 4 nodes in a job
- Also, overall cluster utilization statistics



# Terminal Access

- Naïve to think we can eliminate the command line
- But, we simplify access
  - Zero install
  - Single sign-on (almost, they do have to reauthenticate)
  - Firewall friendly



```
Control keys... websvcs06.osc.edu (Anyterm-1.1.29) Green background :
For questions or assistance, contact
oschelp@osc.edu<mailto:oschelp@osc.edu>.
*****

DOWNTIME ON FEBRUARY 26, 2013: The Ohio Supercomputer Center has
scheduled downtime for all HPC systems on Tuesday, February 26th from
7AM until 5PM. The downtime will affect the Glenn cluster, Oakley
cluster, web portals, and HPC file servers. Login services and access
to storage will not be available during this time.

CHARGING ALGORITHM CHANGE: OSC will adopt a new charging algorithm for
parallel jobs on Oakley and Glenn effective February 26, 2013.
Parallel jobs requesting partial nodes will be charged at a higher
rate after this date. Note that for academic users all charges are in
terms of RUs charged against your allocation.

For more information about the downtime and the charging change, and
to stay up to date on system notices, please visit http://osc.edu/n or
follow @HPCNotices on Twitter.

2013-01-30: As of January 30th, 2013, all parallel Fluent jobs (jobs
that use the "ansyspar" license) will be restricted to the Glenn
system to solve a problem with the scheduler. For more information
please see http://www.osc.edu/supercomputing/software/fluent
*****

As of 2013 Mar 4 04:02 userid dhudak on /nfs/04 used 1GB of quota 500GB and 20 files of quota 1000000 files
As of 2013 Mar 4 04:02 userid dhudak on /nfs/07 used 109GB of quota 500GB and 557205 files of quota 1000000 files
As of 2013 Mar 4 04:02 project/group G-2959 on /nfs/proj01 used 7GB of quota 0GB and 104389 files of quota 0 files
As of 2013 Mar 4 04:02 project/group G-2959 on /nfs/proj02 used 1GB of quota 200GB and 8 files of quota 1000000 files
As of 2013 Mar 4 04:02 project/group G-2959 on /nfs/proj09 used 468GB of quota 2861GB and 878266 files of quota 200000
0 files
As of 2013 Mar 4 04:02 project/group G-2959 on /nfs/proj11 used 110GB of quota 2861GB and 88414 files of quota 2000000
files

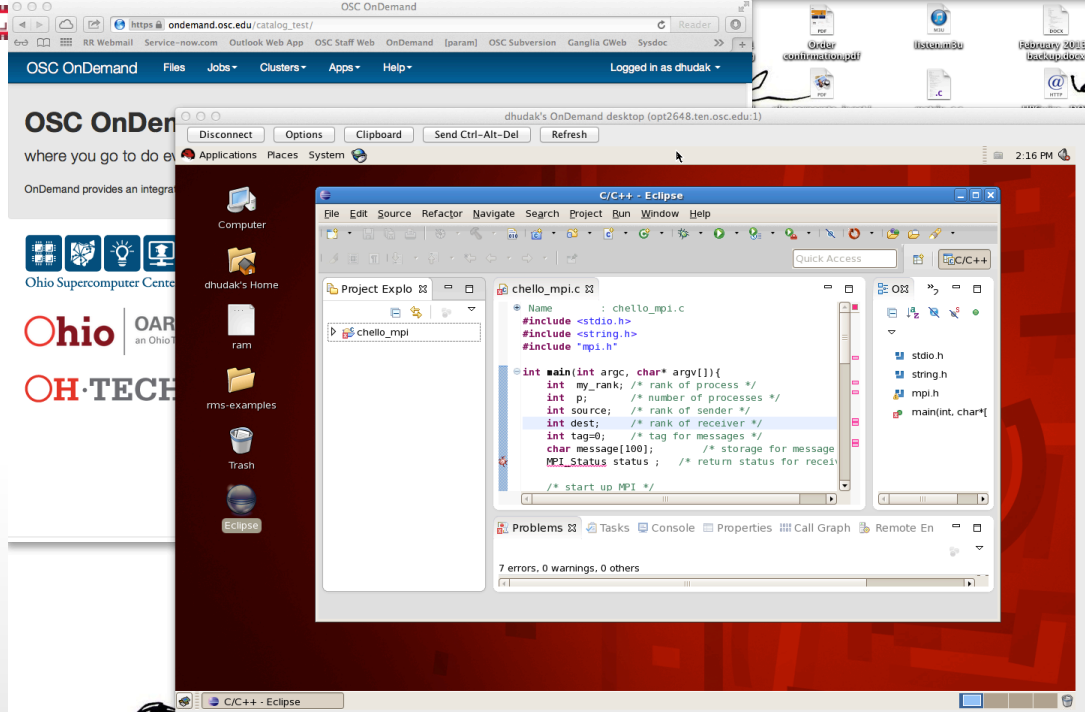
0;-dhudak@oakley02 1001%>
```

- AnyTerm open-source product
- HTML5 terminal emulator
- Automatically log users into login nodes for clusters



# VNC Desktop

- Remote visualization
  - Launching VNC server
  - Launching VNC client
  - Configuring client connection to sever
- We use TurboVNC
- JNLP client (default) or native client

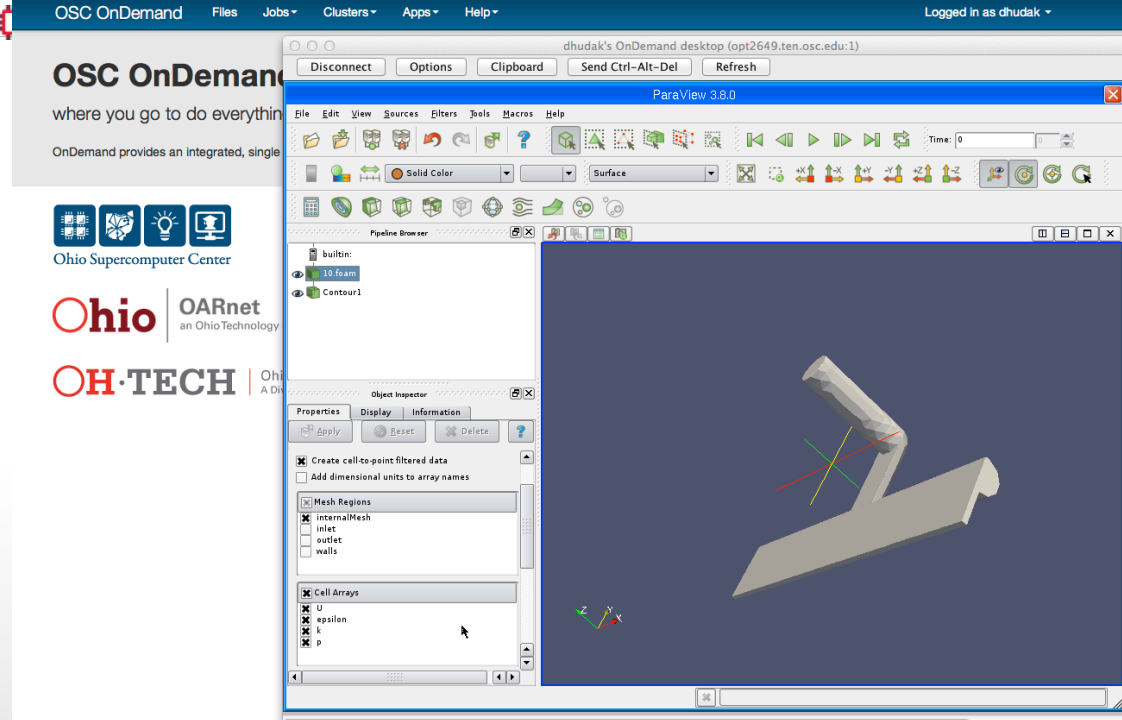


- One click automatically
  - Launches VNC server with 1-time password
  - Configures VNC client with connection information
  - Downloads JNLP client (or connection information) to user



# VNC Applications

- Simplified user experience
  - Abaqus
  - Ansys
  - Comsol
  - ParaView



- Launch VNC connection displaying a single application
- Leveraged in a number of web apps



# Web Applications

- Initial set of apps for industry partners
  - Fluid flow through manifold
  - Thermal dissipation through heat sink
  - Airflow around truck with options (caps and boat tails)
- Common app template (called PUDL for “per-user Drupal”)

The screenshot displays the OSC OnDemand web interface. The top navigation bar includes 'OSC OnDemand', 'Files', 'Jobs', 'Clusters', 'Apps', and 'Help'. The user is logged in as 'dhudak'. The main content area shows a 'CPU Report' for job 'opt2404.ten.osc.edu CPU last custom' and a 'Load Report' for 'opt2404.ten.osc.edu Load last custom'. The CPU Report includes a stacked bar chart and a table of system activity statistics. The Load Report includes a line graph showing system load over time. A dropdown menu is open, listing various web applications under 'WEB APPS' and 'VNC APPS'.

Category	App Name
WEB APPS	Data Shuttle
WEB APPS	Heat Sink Predictor
WEB APPS	Manifold Flow Predictor
WEB APPS	Container Fill Portal
WEB APPS	Truck Add-On Predictor
WEB APPS	Truck Add-On Predictor (dev version)
WEB APPS	Weld Predictor
WEB APPS	GE Weld Predictor
VNC APPS	Glenn Desktop
VNC APPS	Oakley Desktop
VNC APPS	Abaqus
VNC APPS	Ansys
VNC APPS	COMSOL
VNC APPS	Paraview

- Web app workflow
  - Create new job
  - Customize job
  - Submit job
  - Monitor job
  - Visualize results





# OnDemand Implementation

- Traditional cluster architecture
  - Compute nodes, allocated by scheduler
  - Login nodes for interactive use
- We dedicated additional hardware
  - Web nodes run PUWS
  - Visualization nodes run VNC apps

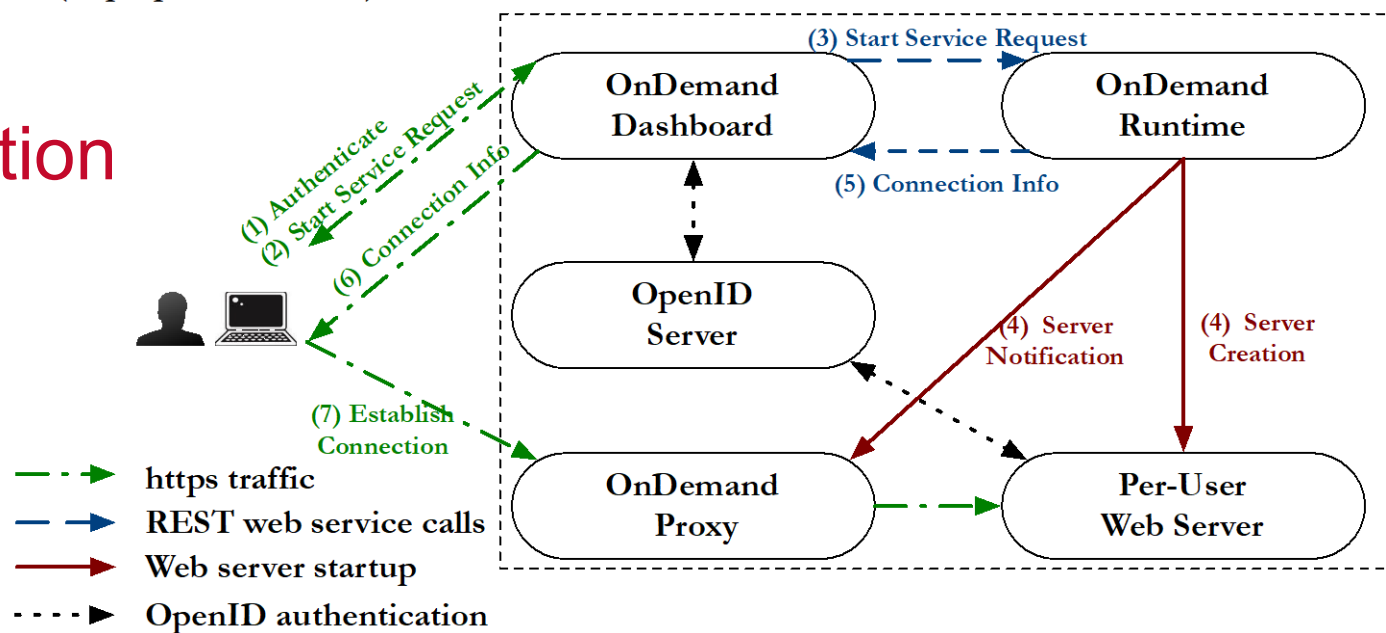


# OnDemand Implementation

- Dashboard
- Runtime
- Proxy
- OpenID

User  
(Laptop and Browser)

HPC Center

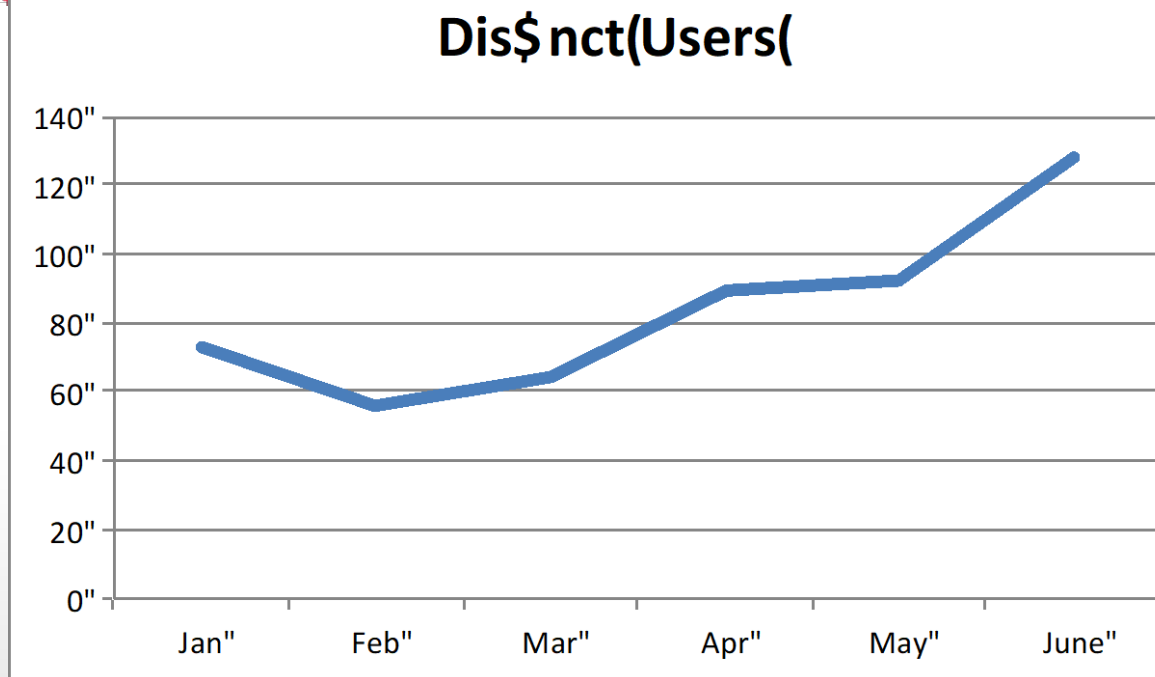


- Seven-step process for PUWS creation
- Proxy enables “firewall friendly” external traffic
  - Proxy told high-number port for PUWS
  - Traffic from user is on https port (443)
  - Proxy makes requests on appropriate high-number port
- PUWS validates OpenID cookie provided by browser with OpenID server
  - Only that authenticated user can access the PUWS



## OnDemand Usage

- “Soft launch” for OnDemand in January 2013
- Friendly users
- Training classes
- Web app users
- Visualization



- General announcement to user group in March
- Steady increase in interested users






## OnDemand Usage

- Distinct users up from 60 to 120 (2X)
- App launches up from 500 to 3000 (6X)
- Indicates users are doing more with OnDemand


## Total App Launches



## OnDemand Usage

- Four app categories shown
  - File
  - Job
  - Terminal
  - VNC
- Not enough web apps to show

### App#Launches#by#Category#



- Job app usage entirely in monitoring
  - Very popular with our user support team



## Conclusions and Future Work

- Product is launched and user acceptance is gratifying
- Many open issues/opportunities still remain
  - Authentication: users would like to use shibboleth authentication from their home institution
  - Authorization: show only user's apps
  - PUDL web app template
    - Authentication and user separation are provided at system level
    - Drupal is overkill for our web apps
    - Smaller, simpler web app kit in design
  - OnDemand runtime portability
    - Documentation and links to allow centers to set up their own



# Conclusions and Future Work

- Many open issues/opportunities still remain
  - File transfer: limited by apache and http
    - Want to integrate sftp and Globus Online options
  - Job management
    - “My Jobs” to have all user’s jobs, not just ones created by app
    - Show performance statistics for completed jobs
    - Git for job templates and sharing user-created templates
  - Terminal does not pass all VT100 tests
    - For example, screen doesn’t work
    - Chrome’s ssh web app does, though and its nice
  - VNC solutions not browser-based, not firewall friendly
    - Interested in HTML5 VNC client, like Guacamole

