

# BLUE WATERS

SUSTAINED PETASCALE COMPUTING

May 11, 2017

## Application-Level Regression Testing Framework using Jenkins

**Timothy A. Bouvet, NCSA, University of Illinois**  
**Reuben D. Budiardja, ORNL, Oak Ridge, Tennessee**  
**Galen W. Arnold, NCSA, University of Illinois**  
**Gregory H. Bauer, NCSA, University of Illinois**



GREAT LAKES CONSORTIUM  
FOR PETASCALE COMPUTATION

CRAY®

## Talk Outline

- Introduction and Motivation
- Implementation & Deployment
  - Basic Configuration
  - Connecting to multiple HPC systems
  - Security consideration, authentication, authorization
- Application Test Structure
- Actual Use Cases

## Introduction & Motivation

- The complexities of monitoring an HPC system:
  - Layers and versions of software stack
  - System-level configurations
  - I/O components (memory, filesystem)
  - Diversity in user need and usage
  - Variable network traffic
- Performance regression may occur over time
  - requires historical data to easily detect
- **Application-level regression testing**

Aggregate  
impact?

## Goals

- Monitor usability and performance
- User-level experience
  - no privilege access, minimum ‘specialized’ tests
- Early detection of regressions
  - Tests are run during production period
- Automated and consistent test solution
- Store historical data
- **Solution: Jenkins Automation Server**

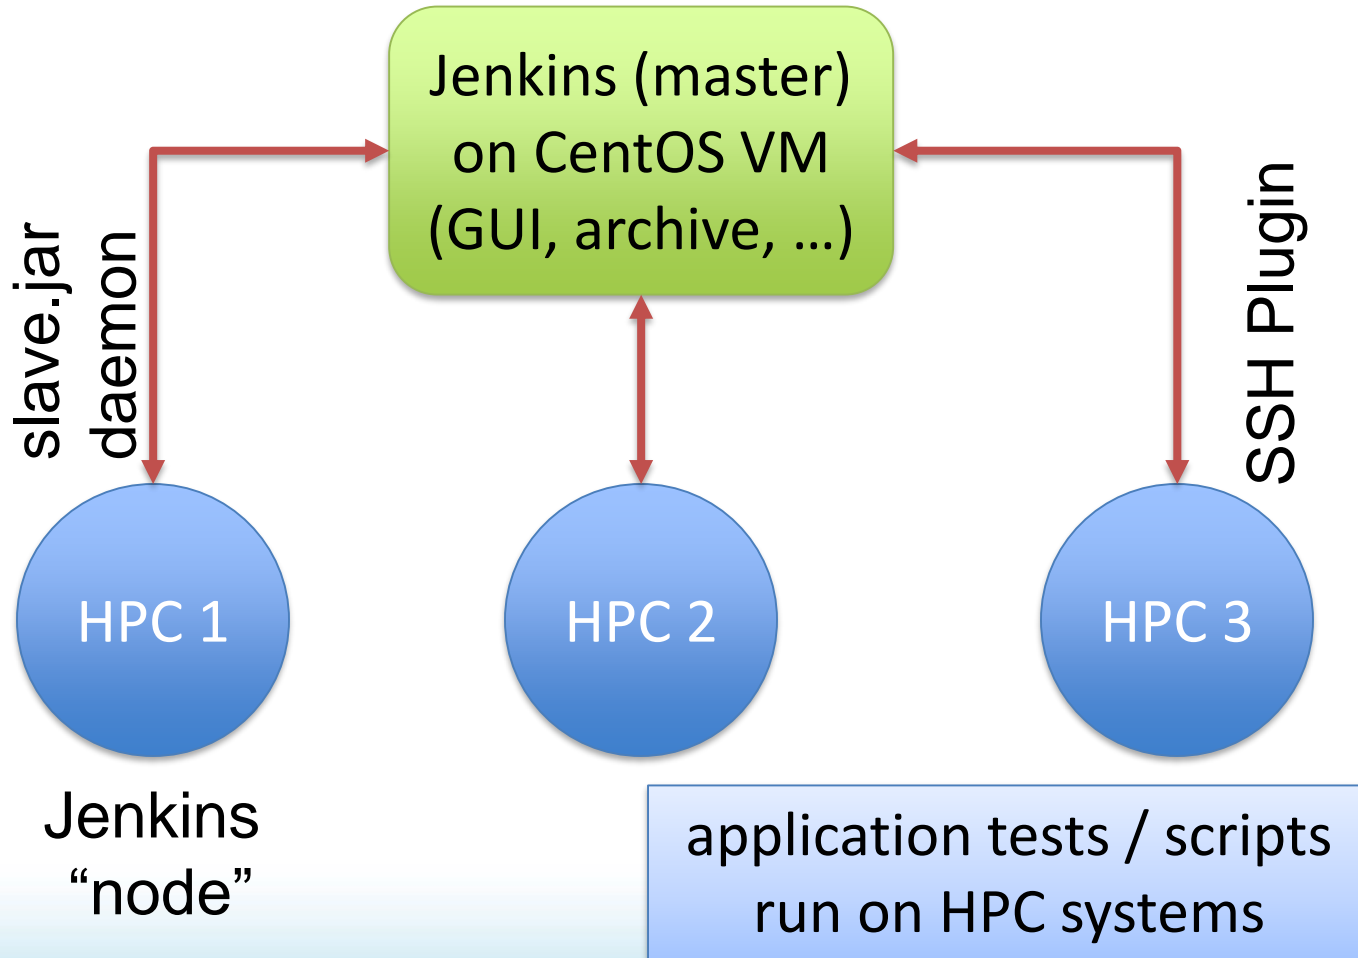
# What is Jenkins?



- An “open source automation server”
- Automate user-level tasks
  - Execute arbitrary commands
  - Build software & submit jobs
  - Archive test results
  - Notify (email, text, ...) on error
- “Never sleeps or takes vacation”
- Chosen for
  - versatility
  - multitudes of plugins (plot, ssh, ... )
  - large community support

# Implementation & Deployment

# Basic Configuration



# Accessing (Remote) HPC Test Systems

## Jenkins “Node”

- Requires running “slave.jar” daemon
- Daemon connects back to master on specific port
- Remote execution appears “local”
  - no need to copy files / data
  - can access lustre, /home as if they are local resources

**NICS**

## With SSH Plugin

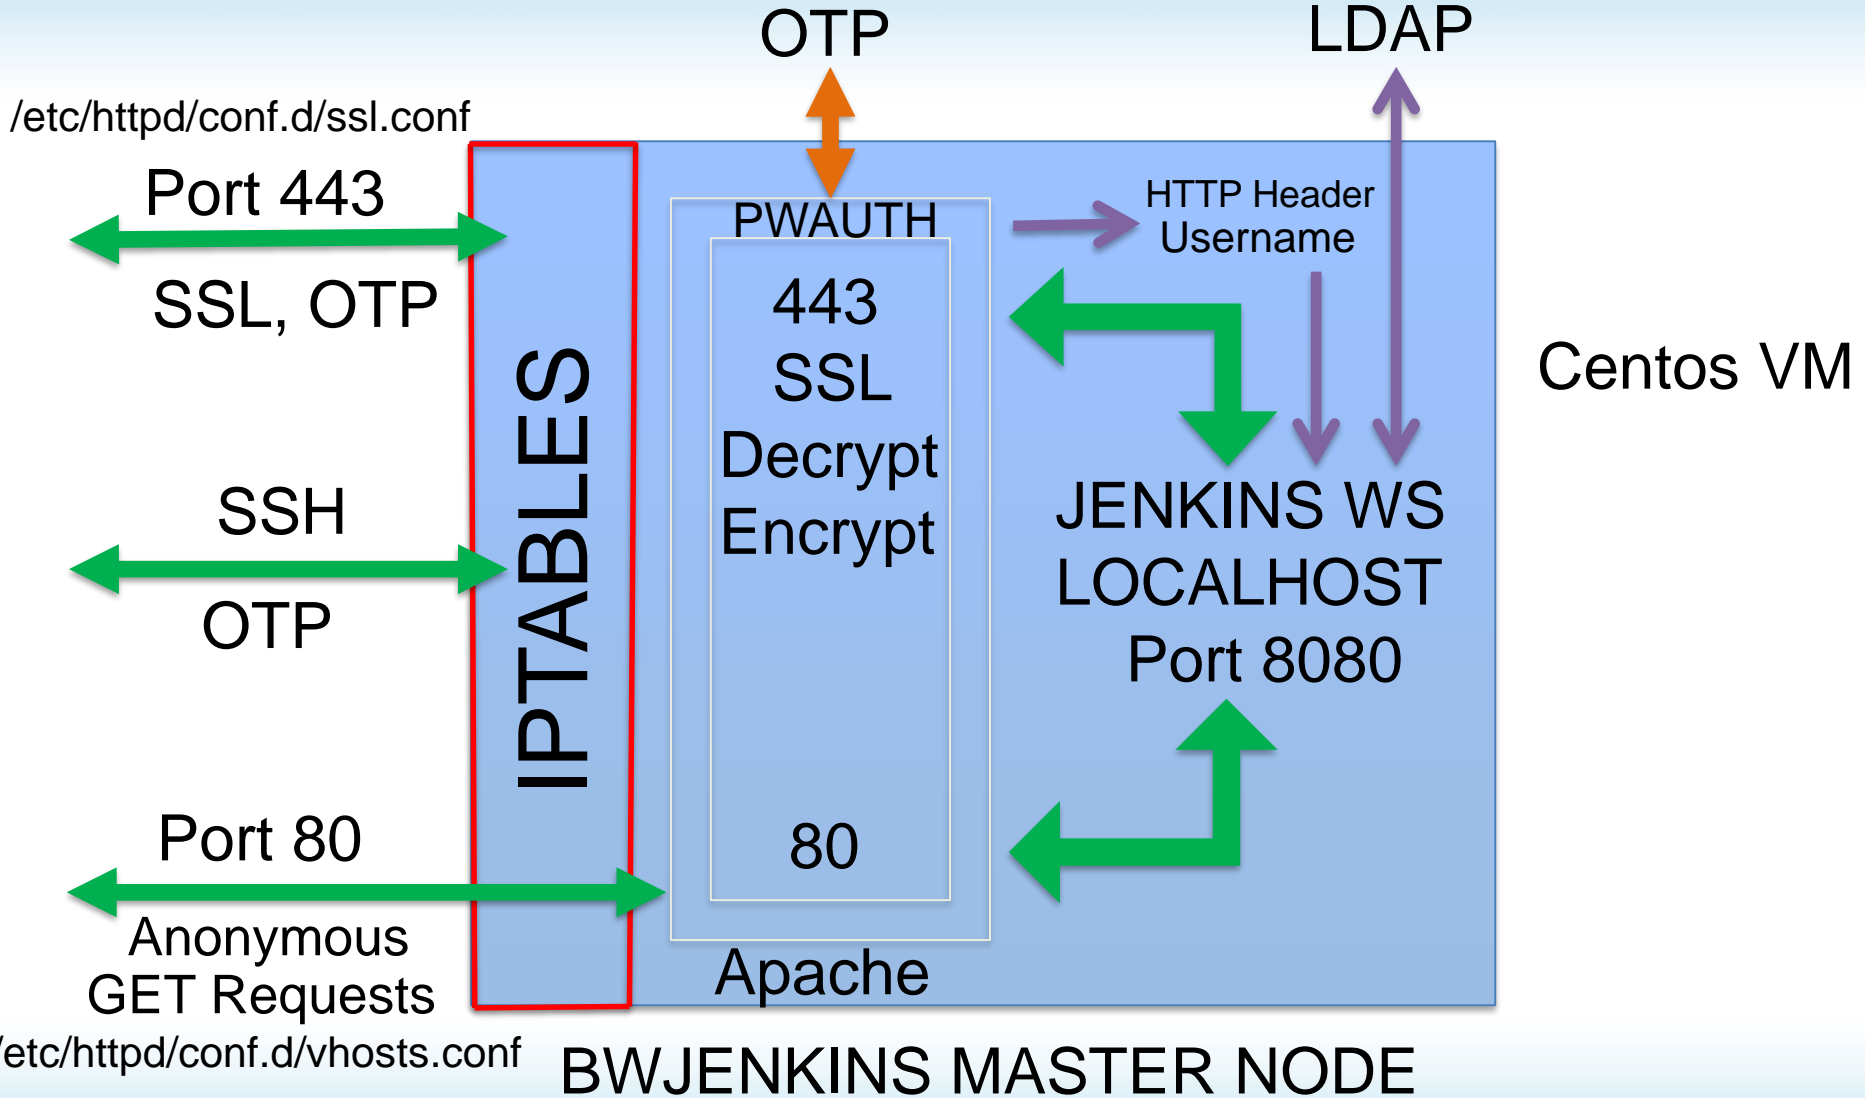
- Need to set up passwordless login with SSH keys
- Closer resemblance to user experience
- Have to manually manage data movement (file copying, getting results back) to

**NCSA**



## Security Considerations

- Jenkins allows execution of arbitrary command
  - use Apache HTTPD server in front of Jenkins', pass requests via mod\_proxy
  - allow authenticated access only from institutional IPs via firewall rules
  - use RSA OTP for authentication via pwauth + custom session management (details in paper)
- Use LDAP for authorization with Jenkins “matrix-based” security



# Jenkins Global Security

## Authorization

- Anyone can do anything
- Legacy mode
- Logged-in users can do anything
- Matrix-based security



User/group	Overall	Credentials	Agent			Job			Run	View	SCM																						
	Administer	ConfigureUpdateCenter	Read	UploadPlugins	ManageDomains	Create	Delete	Update	View	Build	Configure	Connect	Create	Delete	Disconnect	Provision	Build	Cancel	Configure	Create	Delete	Discover	Read	Workspace	Delete	Update	Configure	Create	Delete	Read	Tag		
Sub_UTK	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>		
arnoldg	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>		
bw_admin	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>		
bw_seas	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>		
rbudiard	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>		
Anonymous	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>		

# Application Test Structure

## An Application Test: Jenkins Project

- Name and description - what is tested, frequency, resources
  - Limit number of “builds” kept - conserve server storage
  - Source code management - subversion or git
  - Build trigger - manual or automated schedule
  - Build commands – shell commands to execute test
  - Post-build Actions – generate plots, notifications on failure
- Project can be easily copied from an existing one*

# Admin Gui Home Page

- New Item
- People ← **Anonymous**
- Build History ← **Only View**
- Edit View
- Delete View
- Log In ←
- Manage Jenkins
- My Views
- Credentials

**Build Queue** -

No builds in the queue.

**Build Executor Status** -

**master**

- 1 Idle
- 2 Idle
- 3 Idle
- 4 Idle
- 5 Idle
- 6 Idle
- 7 Idle
- 8 **ior-h2ologin4-darshan** #197
- 9 **MILC-BW** #705
- 10 Idle

Add JYC tests that are considered done and production-ready here with "Edit View".

[edit description](#)

S	W	Name ↓	Last Success	Last Failure	Last Duration	
		<a href="#">cray-hdf5-parallel-jyc</a>	1 day 18 hr - <a href="#">#48</a>	5 mo 19 days - <a href="#">#21</a>	10 min	
		<a href="#">cuda-jyc</a>	7 hr 45 min - <a href="#">#442</a>	N/A	10 min	
		<a href="#">HDF5Benchmarks</a>	19 min - <a href="#">#2095</a>	1 day 1 hr - <a href="#">#2085</a>	58 sec	
		<a href="#">IOR-jyc</a>	8 hr 49 min - <a href="#">#613</a>	N/A	11 min	
		<a href="#">JobLaunch-JYC</a> <span style="color: green;">←</span> <b>Select Project</b>	50 min - <a href="#">#7394</a>	1 day 2 hr - <a href="#">#7367</a>	1 min 34 sec	
		<a href="#">LAMMPS</a>	1 day 0 hr - <a href="#">#2086</a>	1 day 1 hr - <a href="#">#2085</a>	46 min	
		<a href="#">Lustre Check Ost JYC</a>	58 min - <a href="#">#6992</a>	N/A	3.9 sec	<span style="color: red;">←</span> <b>Schedule a Build</b>
		<a href="#">mdtest-jyc</a>	1 hr 58 min - <a href="#">#529</a>	17 days - <a href="#">#496</a>	10 min	
		<a href="#">MILC</a>	4 hr 45 min - <a href="#">#1254</a>	14 days - <a href="#">#1250</a>	31 min	
		<a href="#">NAMD</a> <span style="color: green;">←</span> <b>Pass</b>	2 hr 44 min - <a href="#">#2135</a>	14 days - <a href="#">#2126</a>	29 min	
		<a href="#">NWCHEM</a>	22 hr - <a href="#">#1177</a>	13 days - <a href="#">#1173</a>	4 hr 6 min	
		<a href="#">osu_reduce</a>	10 hr - <a href="#">#418</a>	N/A	10 min	
		<a href="#">Qstat JYC</a>	38 min - <a href="#">#8032</a>	N/A	3.9 sec	
		<a href="#">stream-xe-jyc</a> <span style="color: green;">←</span> <b>Disabled</b>	12 hr - <a href="#">#445</a>	26 days - <a href="#">#397</a>	10 min	
		<a href="#">testexternaljob</a>	N/A	N/A	N/A	
		<a href="#">TestSSH-JYC</a> <span style="color: red;">←</span> <b>Fail</b>	2 days 4 hr - <a href="#">#21</a>	1 day 2 hr - <a href="#">#25</a>	6 sec	

IOR-bluewaters

Would last have run at Thursday, May 11, 2017 7:53:34 AM CDT; would next run at Thursday, May 11, 2017 1:53:34 PM CDT.

Poll SCM

### Build Environment

Execute shell script on remote host using ssh

### Build

Execute shell script on remote host using ssh

SSH site

bwjenkins@bw.ncsa.illinois.edu:22

Command

```
module load autoconf
module load automake
~/bin/myenv.sh
rm -rf jenkins-ior
mkdir jenkins-ior
cd jenkins-ior
git clone https://github.com/LLNL/IOR
cd IOR
./bootstrap
./configure --host=x86_64 CC=cc
make -j 4
file src/ior
ls -l src/ior

cp src/ior ~arnoldg/ior/src/ior
#make clean; make
cat -n ~arnoldg/ior/compare_stripe.pbs # there is a lorscript input file for ior in the same directory
MYJOBID= qsub ~arnoldg/ior/compare_stripe.pbs
echo MYJOBID=$MYJOBID
~/bin/wait_for.sh $MYJOBID

echo "=== JOB RAN ON ==="
cat ~/scratch/stripe8/ior-compute-nodes
if [ -f ~/scratch/stripe8/ior-compute-nodes ]
then
: # It's all good
else
echo "ior-compute-nodes not found, pbs job likely hit walltime"
exit 1
```

# Test Script

Save

Apply

# Project View

Jenkins > BlueWaters > JobLaunch-BW >

Back to Dashboard

Status

Changes

Plots Historical Plots

## Project JobLaunch-BW

Job Launcher Test to Compare Wrapped (XALT) Aprun to Real Aprun on

Tests : environment , xalt , aprun

[Recent Changes](#)

### Build History

find  [Build Details](#) x

<a href="#">#7718</a>	Apr 27, 2017 6:57 PM
<a href="#">#7717</a>	Apr 27, 2017 5:57 PM
<a href="#">#7716</a>	Apr 27, 2017 4:57 PM
<a href="#">#7715</a>	Apr 27, 2017 3:57 PM

### Permalinks

- [Last build \(#7718\), 23 min ago](#)
- [Last stable build \(#7718\), 23 min ago](#)
- [Last successful build \(#7718\), 23 min ago](#)
- [Last failed build \(#7677\), 1 day 17 hr ago](#)
- [Last unsuccessful build \(#7677\), 1 day 17 hr ago](#)
- [Last completed build \(#7718\), 23 min ago](#)



# Example Console Log

Jenkins > BlueWaters > JobLaunch-BW > #7718

Back to Project

Status

Changes

**Console Output**

View as plain text

View Build Information

Previous Build

Select

## Console Output

Started by timer  
Building on master in workspace /var/lib/jenkins/workspace/Jc  
Updating <https://github.com/reubendb/SWTools-BWCI/trunk/jobla>  
At revision 22

No changes for <https://github.com/reubendb/SWTools-BWCI/trunk>  
[JobLaunch-BW] \$ /bin/sh -xe /tmp/jenkins125936332362860387.s  
+ rsync -az --delete -e ssh /var/lib/jenkins/workspace/JobLau  
/bwjenkins/tests/

Access by OTP or Two Factor Certificate Authority only.  
Use myproxy-logon -s tfca.ncsa.illinois.edu -p 7512 for gsis  
gsissh or ssh -o PreferredAuthentications=keyboard-interactiv

Blue Waters Admin Team  
executing script:  
SW\_CONFIG="/scratch/system/bwjenkins/sw\_config"  
PATH="/sw/tools/bin:/sw/tools/bin:\$PATH"  
PYTHONPATH="/sw/tools/bin:/sw/tools/bin:\$PYTHONPATH"  
SW\_ROOT="/scratch/system/bwjenkins/tests"

# Use Cases

# Project Failure Email Notification

See <[http://bwjenkins.ncsa.illinois.edu/job/Lustre\\_Check\\_OST\\_BW/7218/display/redirect](http://bwjenkins.ncsa.illinois.edu/job/Lustre_Check_OST_BW/7218/display/redirect)>

-----  
Started by timer

Building on master in workspace <[http://bwjenkins.ncsa.illinois.edu/job/Lustre\\_Check\\_OST\\_BW/ws/](http://bwjenkins.ncsa.illinois.edu/job/Lustre_Check_OST_BW/ws/)>  
executing script:

SW\_WORKDIR="/scratch/system/bwjenkins/sw\_workdir"

JOB\_NAME="Lustre\_Check\_OST\_BW"

**snx11003-OST010c-osc-ffff885fc0cd3c00: check error:  
Resource temporarily unavailable**

[SSH] exit-status: 255

Build step 'Execute shell script on remote host using ssh' marked build as failure

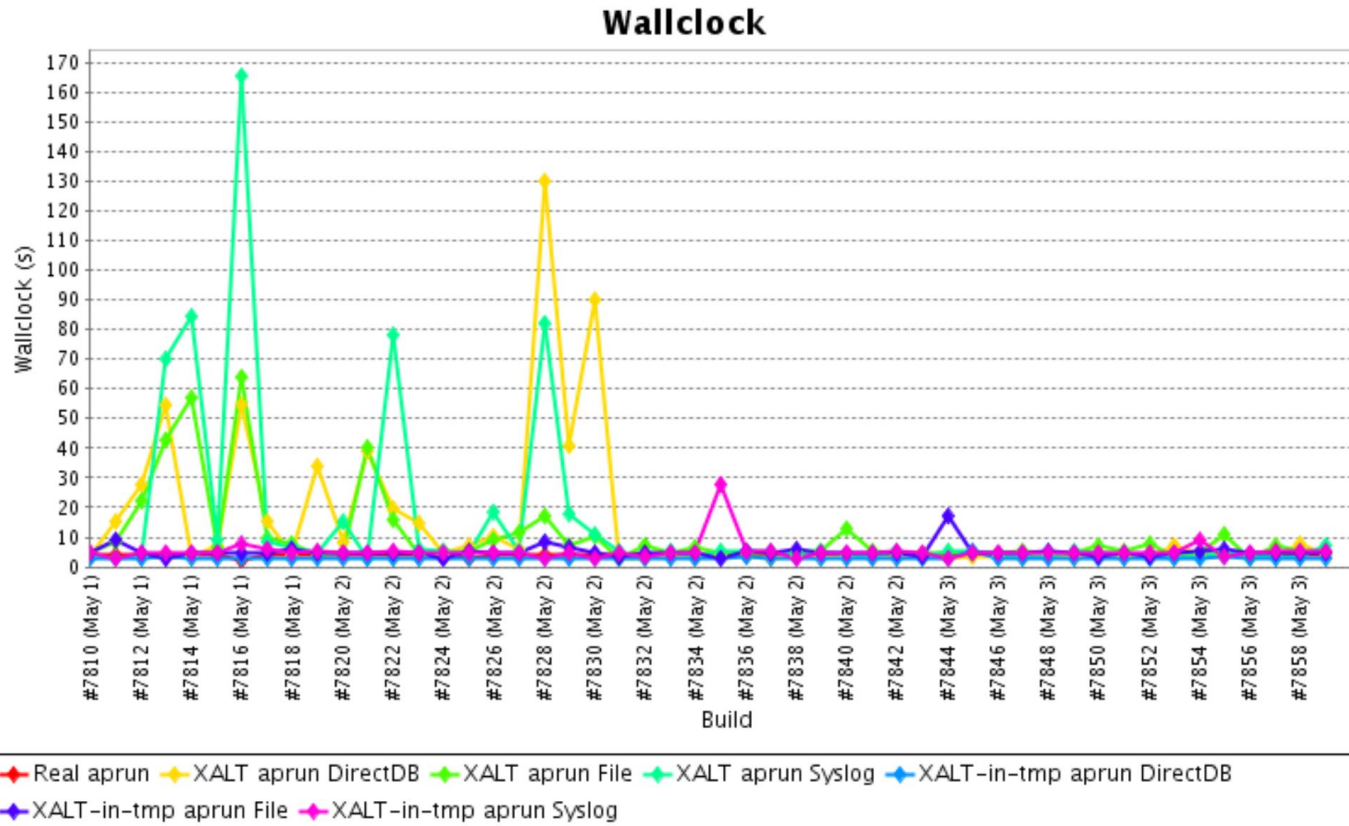
Recording plot data

Saving plot series data from:

<[http://bwjenkins.ncsa.illinois.edu/job/Lustre\\_Check\\_OST\\_BW/ws/bw\\_active](http://bwjenkins.ncsa.illinois.edu/job/Lustre_Check_OST_BW/ws/bw_active)>

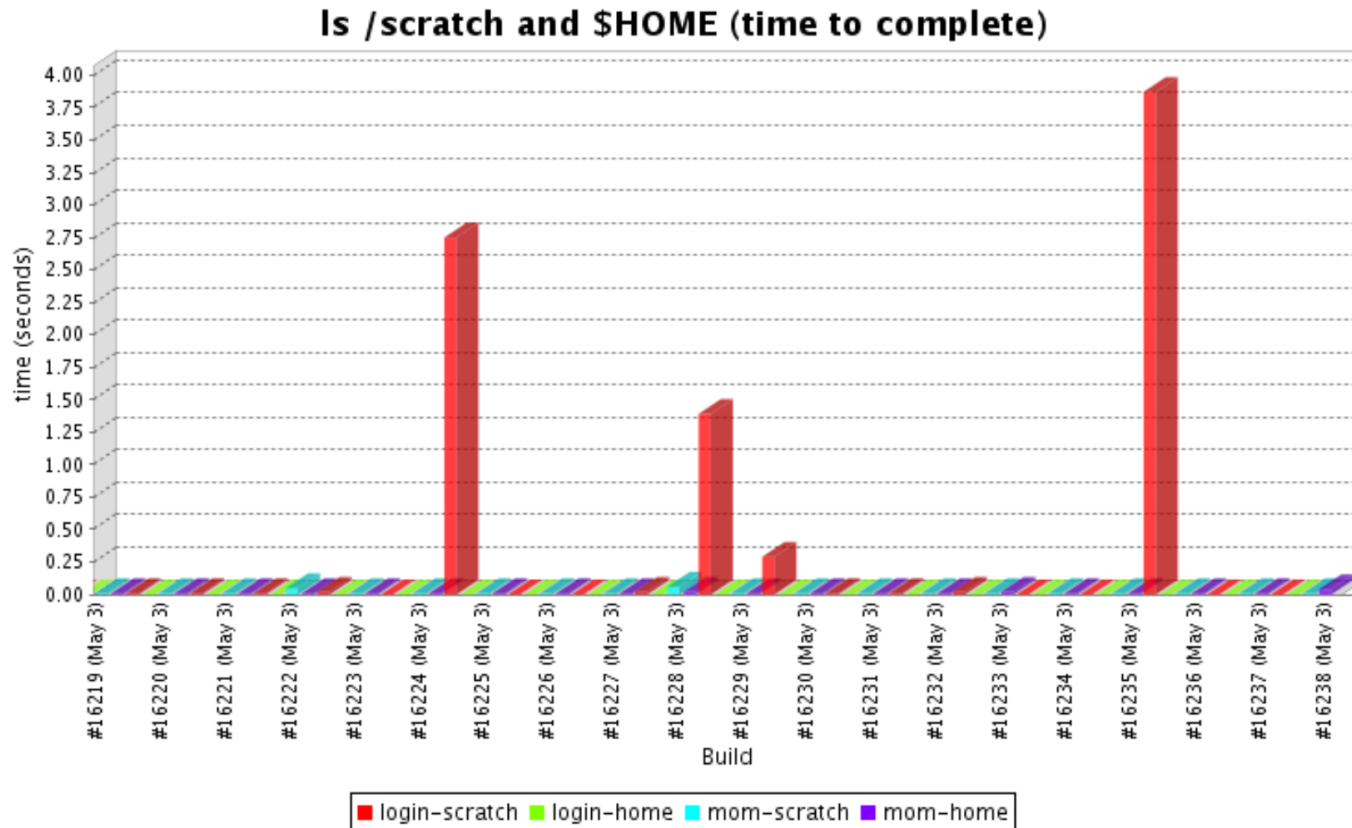
# Slow JobLaunch-BW Plot

[↑ top](#)

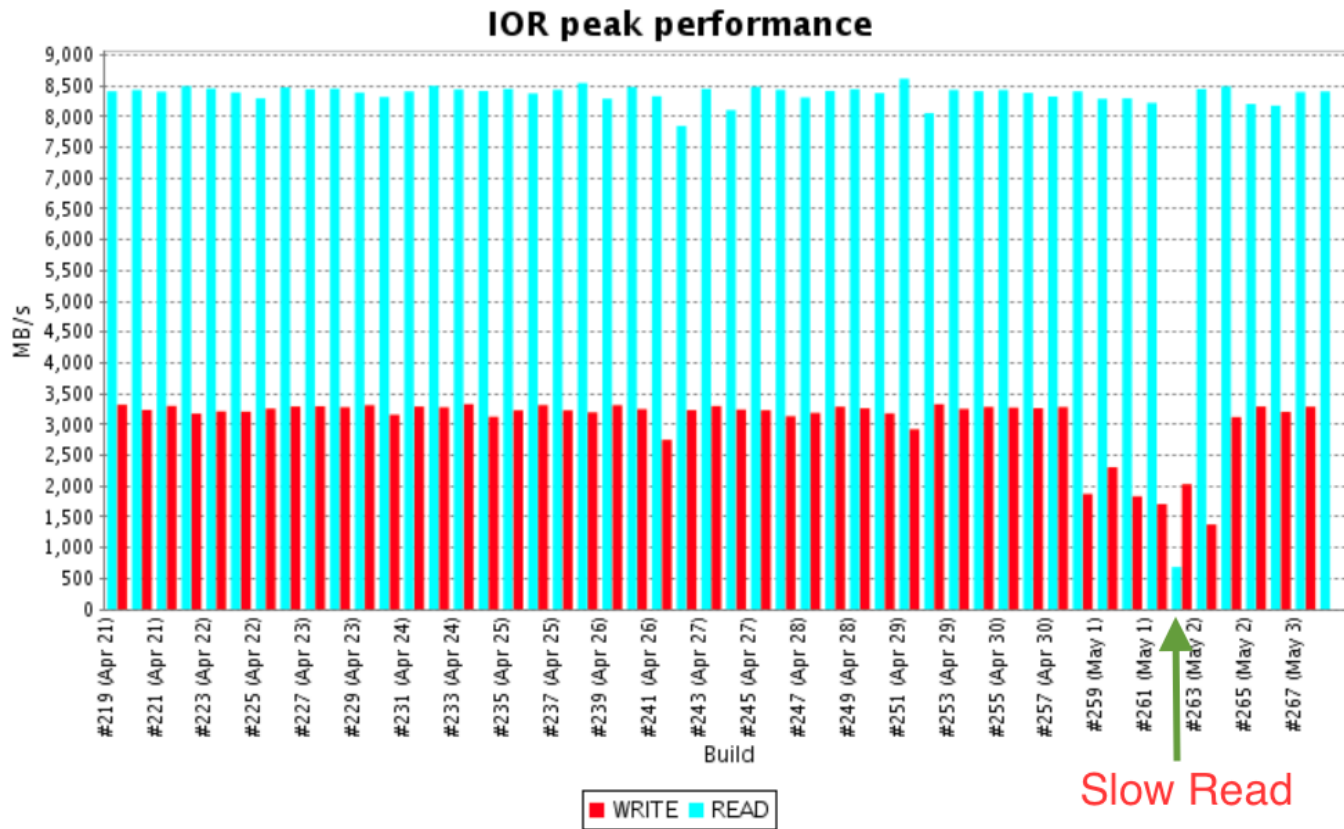


# Slow Lustre Scratch from a login node

[↑ top](#)



# Slow IOR-login Test Example




## IOR-BlueWaters test failures

**Build History** [trend](#)

find

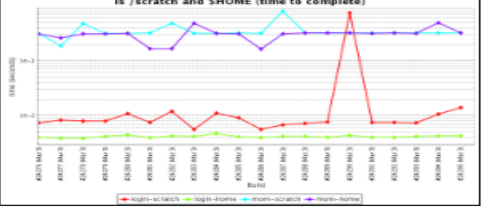

- #1023  
May 4, 2017 1:53 PM
- #1022  
May 4, 2017 7:53 AM
- #1021  
May 4, 2017 1:53 AM
- #1020  
May 3, 2017 7:53 PM
- #1019  
May 3, 2017 1:53 PM
- #1018  
May 3, 2017 7:53 AM
- #1017  
May 3, 2017 1:53 AM
- #1016  
May 2, 2017 7:53 PM
- #1015  
May 2, 2017 1:53 PM
- #1014  
May 2, 2017 7:53 AM
- #1013  
May 2, 2017 2:41 AM
- #1012  
May 1, 2017 7:53 PM



Two Failed Tests that went over 10 minute wallclock.  
These failures indicate an issue with the filesystem.

# Plot Usage : Wiki Filesystem dashboard

Jenkins Monitoring, see also : <http://bwjenkins.ncsa.illinois.edu/view/BlueWaters/>

<p><b>scratch</b></p>	 <p>higher is better. If low only once, build might have used busy OSTs.</p>	 <p>higher is better. Trouble if mdtest and IOR rates are both low.</p>
<p><b>scratch and \$HOME Is test</b></p>	<p>...</p>	 <p>Low is better. This is a good measure of how the filesystem "feels" to users. Times &lt; 1.0s are generally good. These tests are on a 5 minute interval so this graph is approximately the last hour's performance.</p>
<p><b>home</b></p>	<p>No IOR testing: We don't do large block i/o on \$HOME.</p>	 <p>higher is better.</p>



## Conclusion

- Regression testing framework, enabling:
  - Reproducibility of tests
  - Regression testing
  - Rapid reaction ← more tests are added as needed
- Has been deployed at
  - Blue Waters at NCSA
  - NICS
  - Compute and Data Environment for Science (CADES) at ORNL (in progress)
- Future work: integrate with other change managers