



test harness and reporting framework

Shava Smallen
San Diego Supercomputer Center

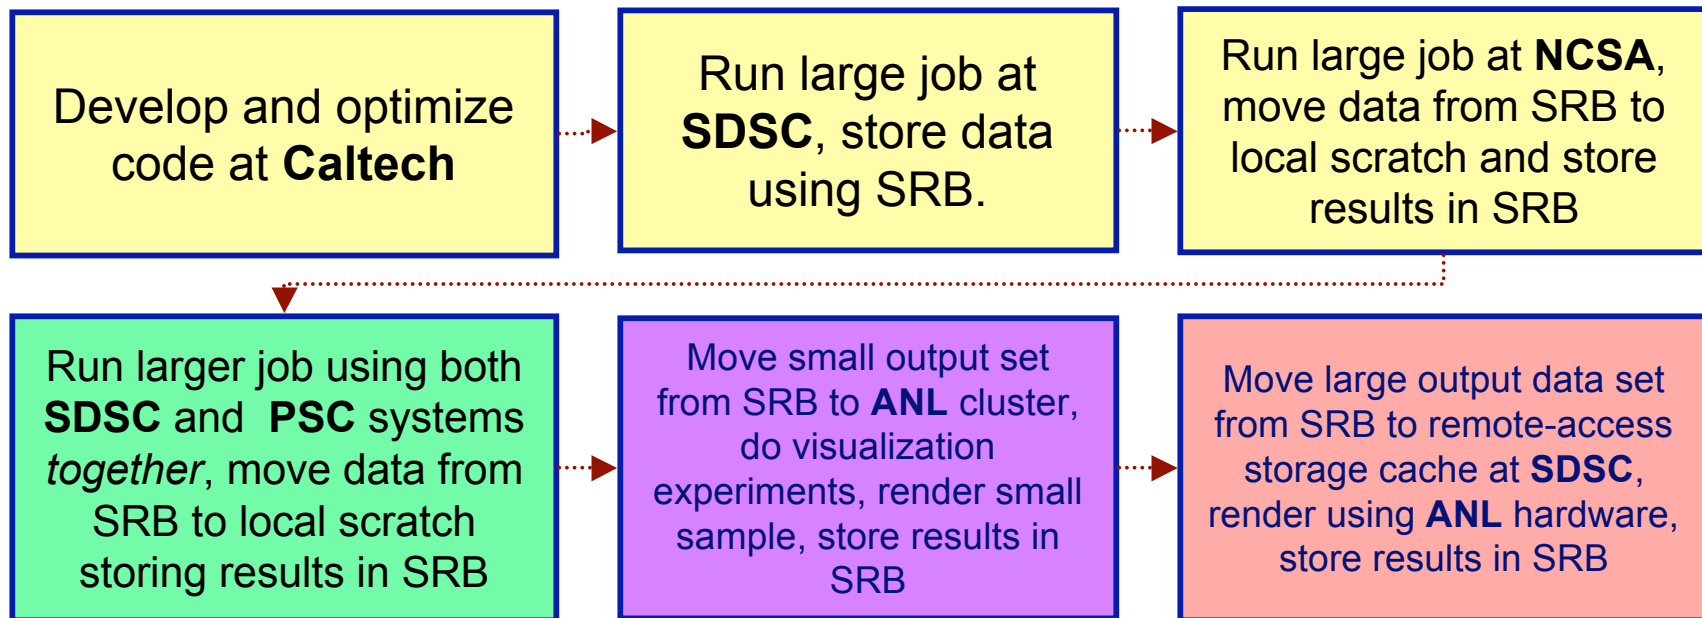
Grid Performance Workshop
6/22/05

Is the Grid Up?

- Can user X run application[s] Y on Grid[s] Z? Access dataset[s] N?
 - Are Grid services the application[s] use available? Compatible versions?
 - Are dataset[s] N accessible to user X? Credentials?
 - Is there sufficient space to store output data?
 - ...
- Community of users (VO)?
- Multiple communities of users?

Testing a Grid

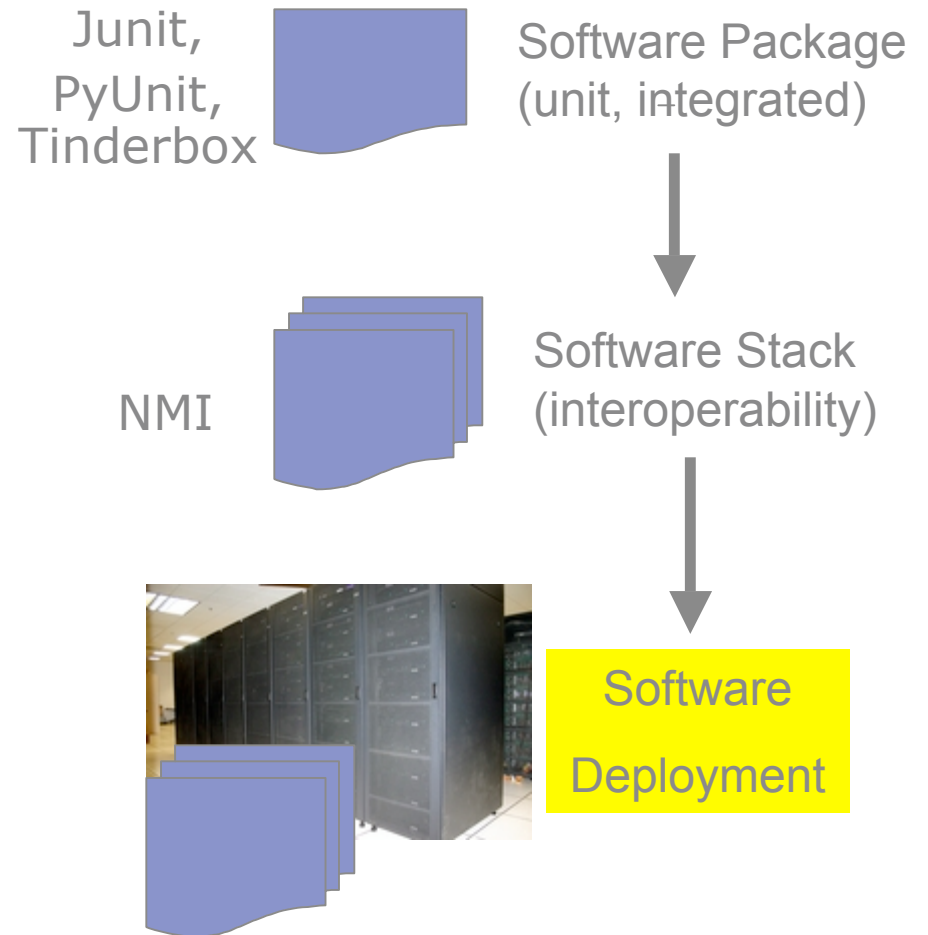
- If you can define “Grid up” in a machine-readable format, you can test it
- User documentation, users, mgmt



Grid up example

What type of testing?

- Deployment testing
 - Automated, continuous checking of Grid services, software, and environment
 - Installed? Running? Configured correctly? Accessible to users? Acceptable performance?
 - E.g., gatekeeper ping or scaled down application

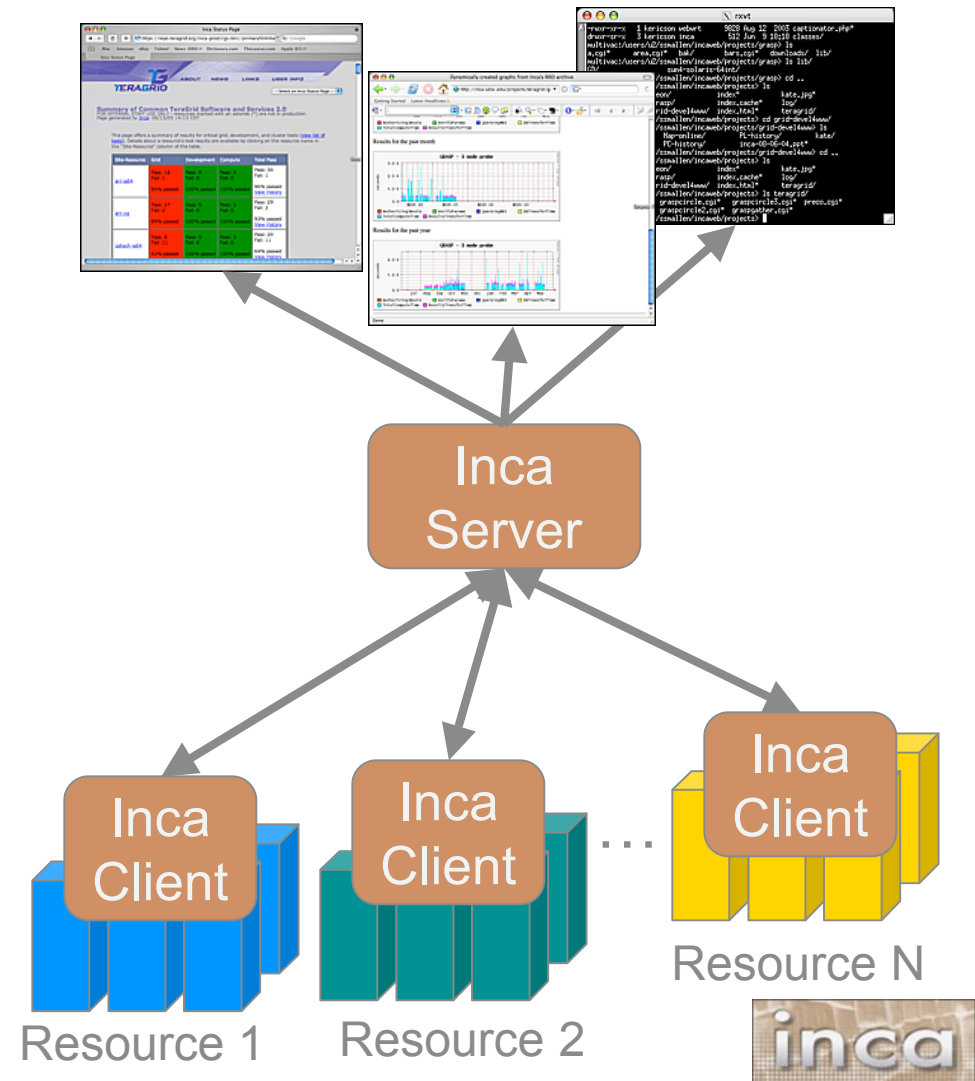


Who tests?

- Grid/VO Management
 - Run from default user account
 - Goal: user-level problems detected & fixed before users notice
 - Results available to users
- User-specific
 - Debug user account/environment issues
 - Advanced usage: feedback tests

Inca

- Framework for the automated testing, benchmarking and monitoring of Grid systems
 - Schedule execution of information gathering scripts (reporters)
 - Collect, archive, publish, and display results



Outline

- ✓ Introduction
- Inca architecture
- Case study: V&V on TeraGrid
- Current and Future Work
- Feedback

Inca Reporters

- Script or executable that outputs XML conforming to Inca specification
- Context of execution is required - important for repeatability
 - What commands were run?
 - What machine?
 - What inputs?
 - What time?
 - What result?
- Communicate more than pass/fail
 - Body XML can be reporter specific - flexibility
 - E.g., package version info (software stack availability)
 - E.g., SRB throughput (unusual drop in SRB performance)
- Users can run it independently of framework

Reporter Execution Framework

- How often should reporters run
 - Boot-time, every hour, every day?
- Modes of execution:
 - One shot mode:
 - boot-time, after a maintenance cycle, user checking their specific setup
 - Continuous mode: cron scheduling
- Data can be queried from a web service and displayed in a web page

Outline

- ✓ Introduction
- ✓ Inca architecture
- Case study: V&V on TeraGrid
- Current and Future Work
- Feedback

TeraGrid

- TeraGrid - an “enabling cyberinfrastructure” for scientific research
 - ANL, Caltech, Indiana Univ., NCSA, ORNL, PSC, Purdue Univ., SDSC, TACC
 - 40+ TF, 1+ PB, 40Gb/s net
- Common TeraGrid Software & Services
 - Common user environment across heterogeneous resources
 - TeraGrid VO service agreement



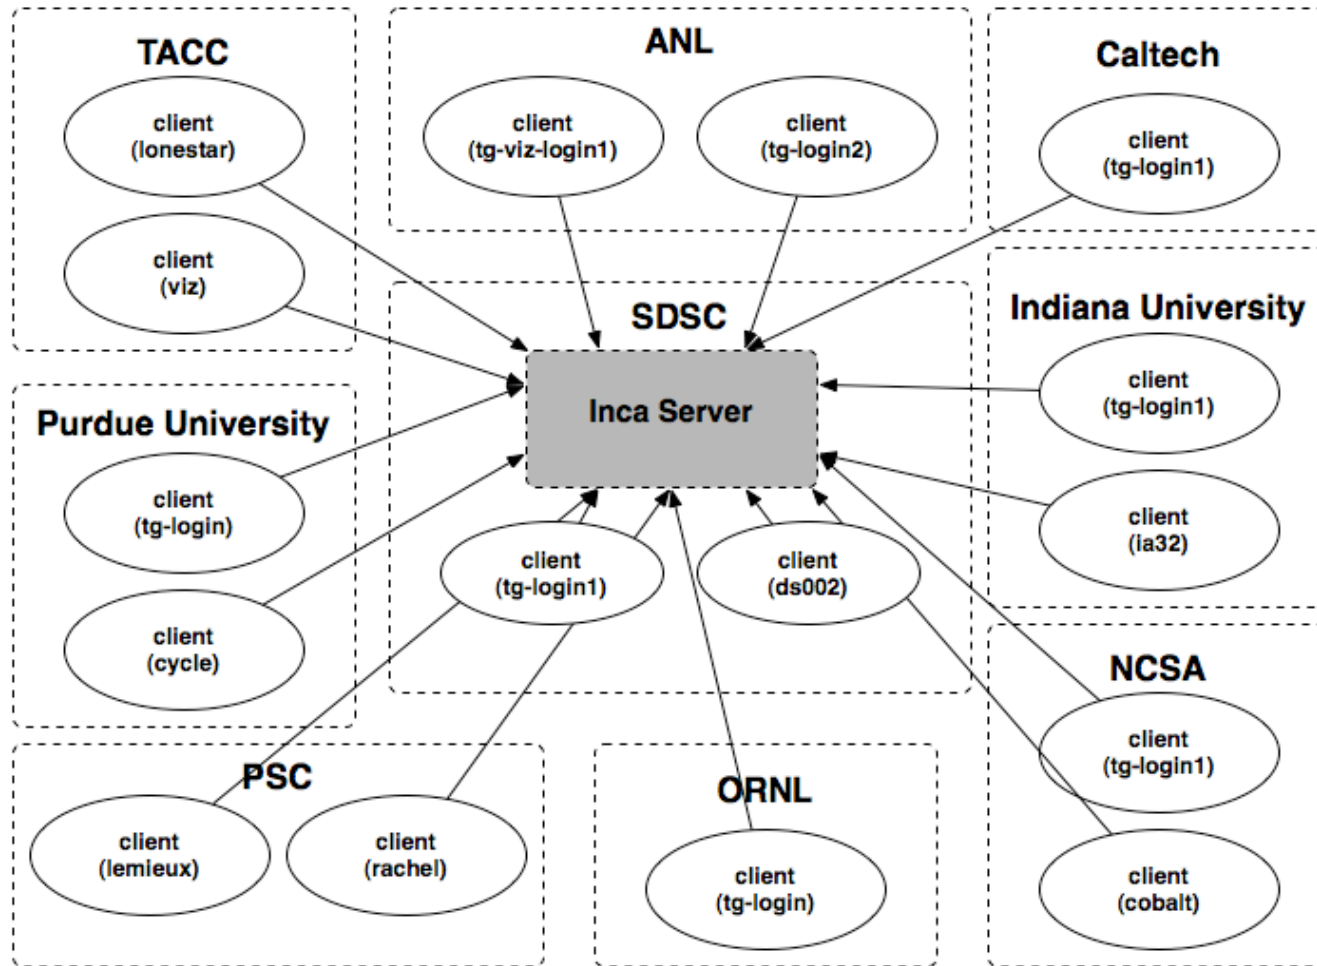
Validation & Verification

- Common software stack:
 - **20 core packages**: Globus, SRB, Condor-G, MPICH-G2, OpenSSH, SoftEnv, etc.
 - **9 viz package/builds**: Chromium, ImageMagick, Mesa, VTK, NetPBM, etc.
 - **21 IA-64/Intel/Linux packages**: glibc, GPFS, PVFS, OpenPBS, intel compilers, etc.
- ➡ 50 version reporters: compatible versions of SW
- ➡ 123 tests/resource: package functionality
 - **Services**: Globus GRAM, GridFTP, MDS, SRB, DB2, MyProxy, OpenSSH
 - **Cross-site**: Globus GRAM, GridFTP, OpenSSH

Validation & Verification (cont.)

- Common user environment
 - \$TG_CLUSTER_SCRATCH,
\$TG_APPS_PREFIX, etc.
 - SoftEnv configuration - manipulate user environment
- ➔ Verify environment vars defined in default environment
- ➔ Verify Softenv keys defined consistently across sites

Inca deployment on TeraGrid



- 9 sites/16 resources
- Run under user account **inca**

Detailed Status Views

[illegible]

Drill-down capability

Reporter details:

reporter name	grid.middleware.globus.unit.gatekeeper <i>(click on reporter name to view reporter script)</i>
description	This test runs globusrun -a [hostname] to check that the gatekeeper at the host is accessible from the local machine
version	1.4
status	production
url	http://www.ncsa.uiuc.edu/People/jbasney/teragrid-setup-test.html

Execution information:

inputs	verbose 1
	help no
	log 3
	host test_hostname
ran at (GMT)	Wed Jun 15 00:13:02 2005
age	27 mins
runs every	1 hour(s)

Reporter system command log:

The following are the *system* commands executed by the reporter. Note that the reporter may execute other actions in between system commands (e.g., change directories). Please click the on reporter name above for the full reporter code.

```
% globusrun -a -r test_hostname 2>&1
```

Host information:

hostname	ran_on_hostname
ipaddr	192.100.00.000
uname	Linux ##### SMP Fri Jun 3 11:44:48 EST 2005 i686 i686 i386 GNU/Linux

Summary Status

TERAGRID [Summary](#) * [CTSS](#) * [SoftEnv](#) ([detail](#)) * [Default User Environment](#)

Summary of Common TeraGrid Software and Services 2.0

Page generated by [Inca](#): 07/13/04 18:39 CDT

This page offers a summary of results for critical grid, development, and cluster tests ([view list of tests](#)). Details about a resource's test results are available by clicking on the resource name in the "Site-Resource" column of the table.

Site-Resource	Grid	Development	Cluster	Total Pass
site1-resource1	Pass: 32 Fail: 1 96% passed	Pass: 23 Fail: 0 100% passed	Pass: 1 Fail: 1 50% passed	Pass: 56 Fail: 2 96% passed
site1-resource2	Pass: 22 Fail: 7 75% passed	Pass: 23 Fail: 0 100% passed	Pass: 1 Fail: 1 50% passed	Pass: 46 Fail: 8 85% passed
site2-resource1	Pass: 1 Fail: 18 5% passed	Pass: 2 Fail: 10 16% passed	n/a	Pass: 3 Fail: 28 9% passed

Expanded View of Errors

site1-resource1

Grid


- globus-2.4.3-intel-r3: [failed: duroc_mpi_helloworld_to_jobmanager-pbs test](#)

All tests passed: 100%

One or more tests failed: < 100%

Tests not applicable to machine or have not yet been ported



Key



History of percentage of tests passed in "Grid" category for a 6 month period

Measuring TeraGrid Performance

- GRASP (Grid Assessment Probes)
 - test and measure performance of basic grid functions
- Pathload [Dovrolis et al]
 - measures dynamic available bandwidth
 - uses efficient and lightweight probes



Lessons learned

- Initially focused on system administrative view
- Moving towards user-centric view
 - File transfer functionality and performance
 - File system availability
 - Job submission
 - SRB performance
 - Interconnect bandwidth
 - Applications: NAMD, AWM

Integr

Base

Are you
•[Data](#)

[YES](#): Are
a file?

[YES](#): Are

1. [Check](#)
[proxy](#)

...

Report

Verifies t

Option

Action:

Run
Run
View

Submit

grid.middleware.globus.unit.proxy

Reporter passed.

Reporter details:

reporter [grid.middleware.globus.unit.proxy](#) (click on
name reporter name to view reporter script)
description Verifies that user has valid proxy; attempts to
create if not
version 1.5
status production
url <http://www.globus.org/security/proxy.html>

Execution information:

inputs help no
log 3
verbose 1
ran at Wed Jun 15 23:31:56 2005
(GMT)

Reporter system command log:

The following are the *system* commands executed by the
reporter. Note that the reporter may execute other actions in
between system commands (e.g., change directories). Please click
the on reporter name above for the full reporter code.

```
% grid-proxy-info -timeleft 2>&1
```

Host information:

hostname tg-login1.sdsc.teragrid.org
ipaddr 198.202.112.33
uname Linux tg-login1 2.4.21.SuSE_286.bef2 #1 SMP Wed
May 4 09:24:24 CDT 2005 ia64 unknown

W

n

area

ow

set of

ters

ers

run

Outline

- ✓ Introduction
- ✓ Inca architecture
- ✓ Case study: V&V on TeraGrid
- Current and Future Work
- Feedback

Inca Today

- Software available at:
<http://inca.sdsc.edu>
- Current version: 0.10.3
- Also available in NMI R7
- Users:



DEISA

GEON



<http://inca.sdsc.edu>



Inca 2.0

- Initial version of Inca focused on basic functionality
- New features:
 - Improved storage & archiving capabilities
 - Scalability - control and data storage
 - Usability - improved installation and configuration control
 - Performance - self-monitoring
 - Security - SSL, proxy delegation
 - Condor integration
- Release in 3-6 months

View Error History

Error:

```
call to globus-url-copy failed: error: the server sent an error
response: 425 425 Can't open data connection.
data_connect_failed() failed: an authentication operation
failed.
```

This error has been detected:

7 times in the past
week
20 in the past
month

[[view graph of past error
occurences](#)]

*Submit information or possible
solution for this error:*


[[search for information submitted
for an error or reporter](#)]

Submit information
or suggestions

Search for
information on
error/reporter

Reporter details:

reporter name	data.transfer.gridftp.unit.copy (click to view reporter script)
description	This test verifies the globus-url-copy and destination. If the source file does not contain a small test file.
version	1.9
status	production
url	http://www.ncsa.uiuc.edu/People/jbasney/teragrid-setup-test.html






View Resource Usage

Execution information:

inputs
verbose 1
dstURL gsiftp://resource11.teragrid.org/~inca.gridftp.test
srcURL file:/tmp/inca.gridftp.test
help no
log 3

ran at (GMT) Thu Jun 16 06:34:03 2005
run time (seconds) 5.24 [[graph run time history](#)]
CPU time (seconds) 3.12 [[graph CPU time history](#)]
memory used (MB) 5 [[graph memory used history](#)]
age 16 hours 7 mins
runs every 24 hour(s)



Reporter system command log:

The following are the *system* commands executed by the reporter. Note that the reporter may execute other actions in between system commands (e.g., change directories). Please click the on reporter name above for the full reporter code.

```
% globus-url-copy file:/tmp/inca.gridftp.test  
gsiftp://resource11.teragrid.org/~inca.gridftp.test 2>&1
```

Host information:

hostname ran.on.hostname
ipaddr 192.000.00.000
uname Linux unamenum #2 SMP Fri Jun 3 11:44:48 EST 2005 i686
i686 i386 GNU/Linux

Summary

- Inca is a framework that provides automated testing, benchmarking, and monitoring
 - Grid-level execution to detect problems and report to system administrators
 - Users can view status pages and compare to problems they see
 - Users can run reporters as themselves to debug account/environment problems
- Currently in-use for TeraGrid V&V, GEON, and others

Outline

- ✓ Introduction
- ✓ Inca architecture
- ✓ Case study: V&V on TeraGrid
- ✓ Current and Future Work
- Feedback

Feedback

- How are you monitoring your Grid infrastructure?
- What do you need to test?
- What diagnostic/debugging tools are available to users?
- Displaying test results to users? In what format? How much detail?

More Information

<http://inca.sdsc.edu>

- Current Inca version: 0.10.3
- New version in 3-6 months
- Email:

ssmallen@sdsc.edu