
Inca User-level Grid Monitoring

Shava Smallen

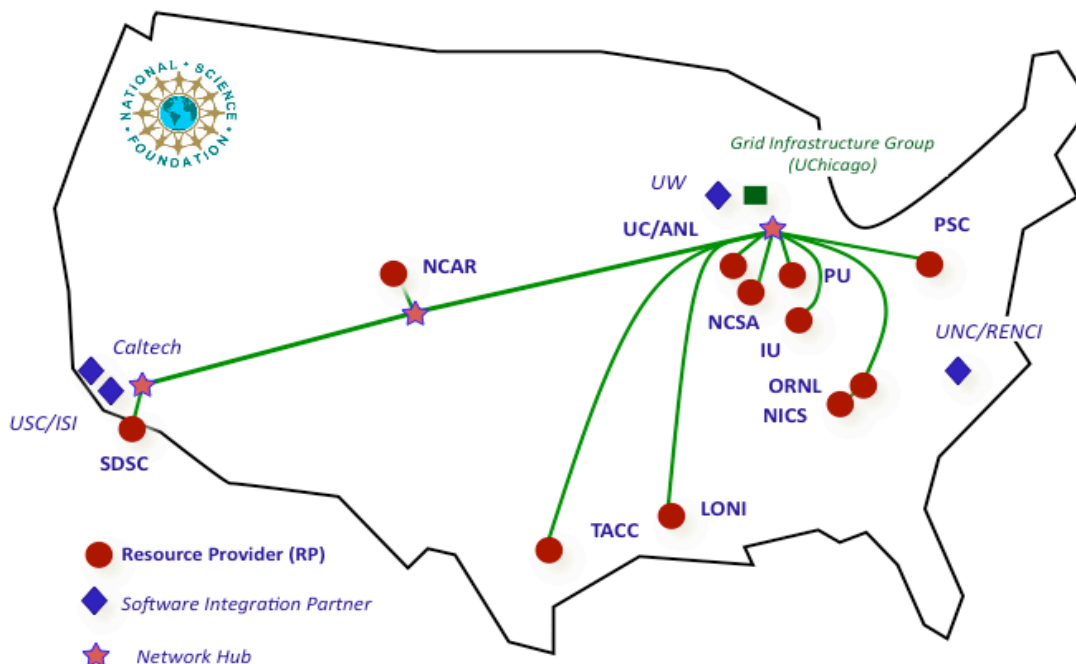
ssmallen@sdsc.edu

SC'09

November 17, 2009

Goal: reliable grid software and services for users

- Over 750 TF
- Over 30 PB of online and archival data storage
- Connected via dedicated multi-Gbps links
- 30-63 software packages and 6-23 services per resource



11 TeraGrid sites, 21 resources




Related Grid monitoring tools

BIG BROTHER™



Nagios®



**Service Availability
Monitoring**

Inca's primary objective: user-level Grid monitoring


User-level grid monitoring

- Runs from a standard user account
- Executes using a standard GSI credential
- Uses tests that are developed and configured based on user documentation
- Automates periodic execution of tests
- Verifies user-accessible Grid access points
- Centrally manages monitoring configuration
- Easily updates and maintains monitoring deployment



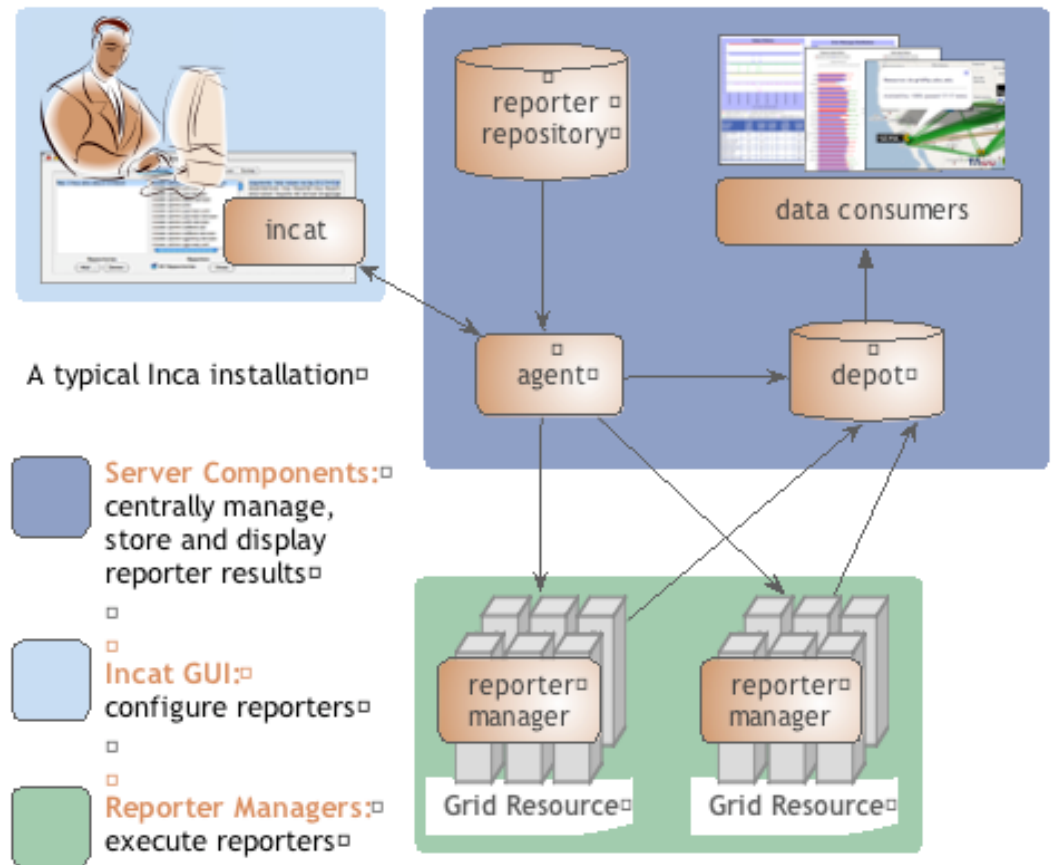
Who benefits from user-level grid monitoring?

- Grid managers
 - Verify requirements are fulfilled by resource providers
 - Identify failure trends
- System administrators
 - Email notification
 - Debugging support
- End users
 - Debug user account/environment issues
 - Advanced users: feedback to Grid/VO



Inca provides user-level grid monitoring

- Stores and archives a wide variety of monitoring results
- Captures context of monitoring result as it is collected
- Eases the writing, deploying, and sharing of new tests or benchmarks
- Flexible and comprehensive web status pages
- Secure



Libraries support common reporter tasks

| Reporter Purpose | Perl Library | Python Library |
|----------------------------|--|---------------------------------------|
| General report | <code>Inca::Reporter</code> | <code>inca.Reporter</code> |
| Software version testing | <code>Inca::Reporter::Version</code> | <code>inca.VersionReporter</code> |
| Software unit testing | <code>Inca::Reporter::SimpleUnit</code> | <code>inca.SimpleUnitReporter</code> |
| Globus unit testing | <code>Inca::Reporter::GlobusUnit</code> | <code>inca.GlobusUnitReporter</code> |
| System performance testing | <code>Inca::Reporter::Performance</code> | <code>inca.PerformanceReporter</code> |

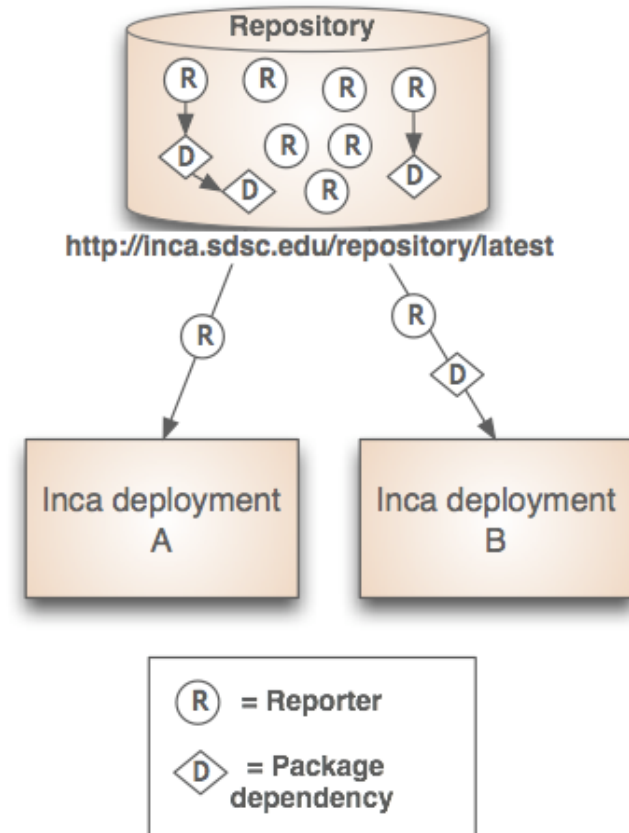
Documentation

<http://inca.sdsc.edu/releases/latest/repdocs/perl.html>

<http://inca.sdsc.edu/releases/latest/repdocs/python.html>

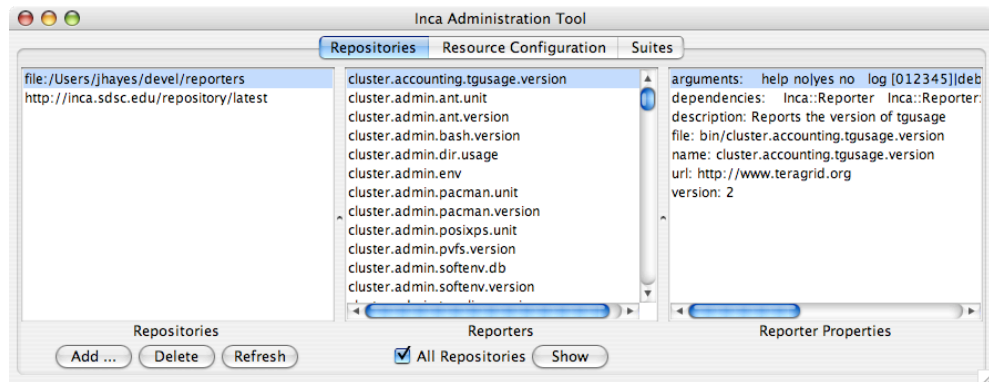
Repositories support sharing

- Collection of reporters available via a URL
- Supports package dependencies
- Packages versioned to allow for automatic updates
- Inca project repository contains 150+ reporters
 - Version, unit test, performance benchmark reporters
 - Grid middleware and tools, compilers, math libraries, data tools, and viz tool



Agent provides centralized configuration and management

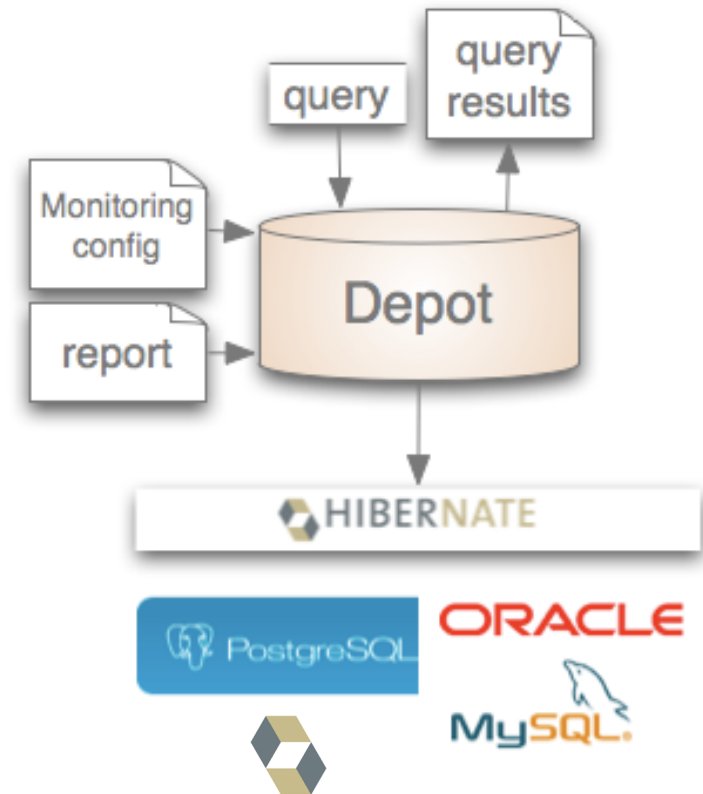
- Implements the configuration specified by Inca administrator
- Stages and launches a reporter manager on each resource - local, SSH, GRAM, WS-GRAM
- Sends package and configuration updates
- Manages proxy information
- Administration via GUI interface (incat)



Screenshot of Inca GUI tool, incat, showing the reporters that are available from a local repository

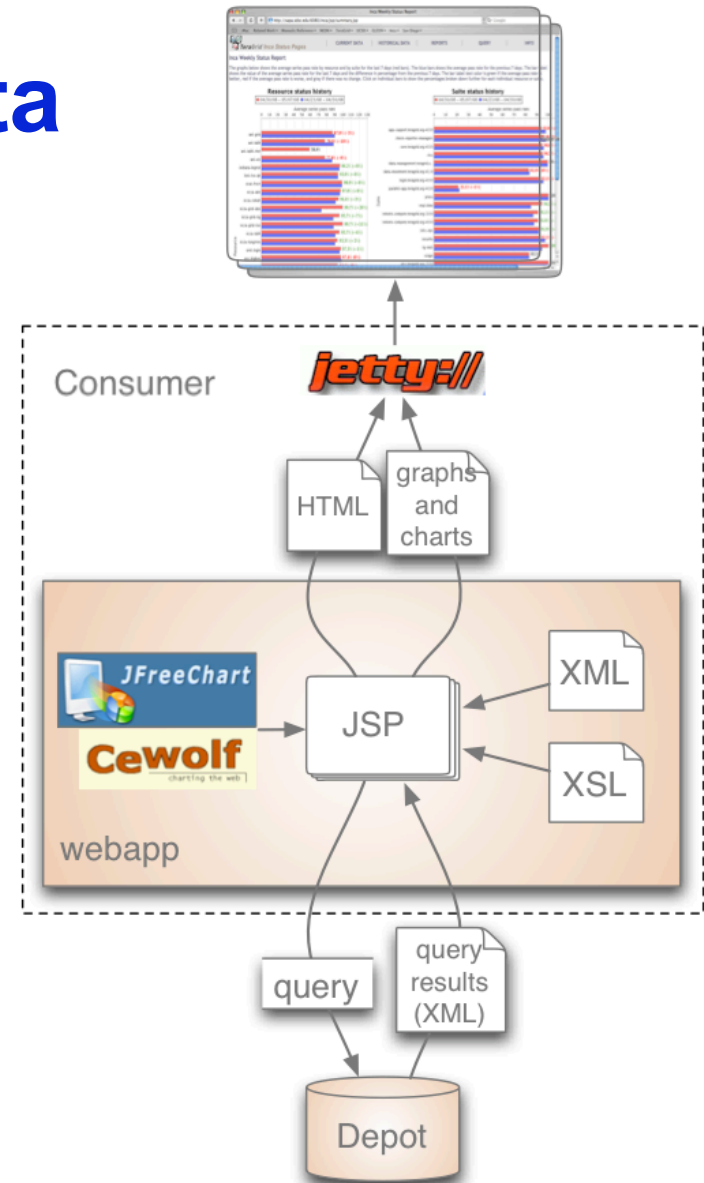
Depot stores and publishes data

- Stores configuration information and monitoring results
- Provides full archiving of reports
- Uses relational database backend via Hibernate
- Supports HQL and predefined queries
- Supports plug-in customization (e.g., email notifications, downtimes)
- Web services - Query data from depot and return as XML



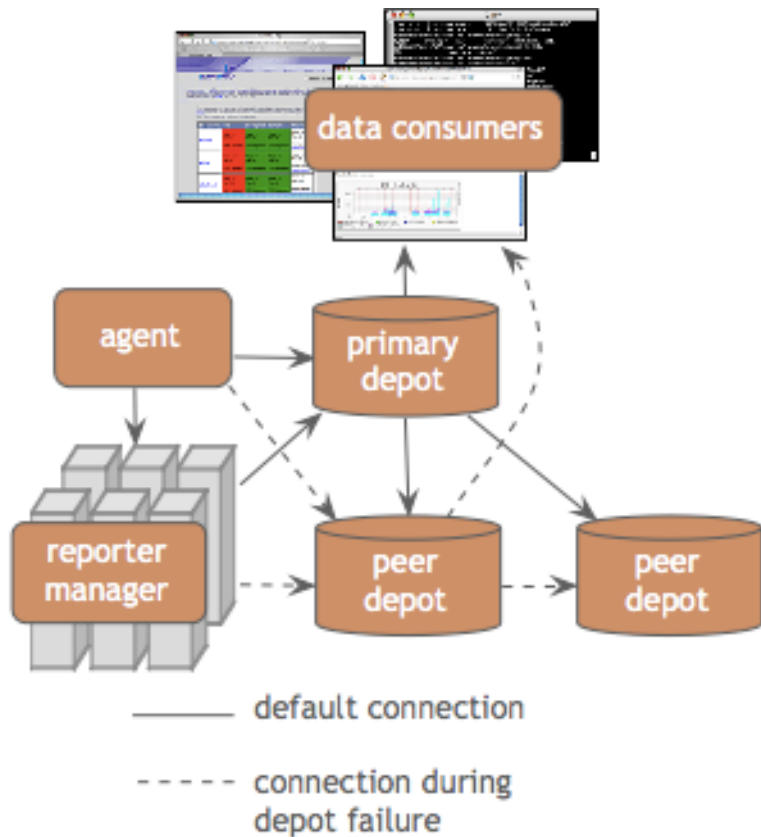
Consumer displays data

- Current and historical views
- Web application packaged with Jetty
- JSP 2.0 pages/tags to query data and format using XSLT
- CeWolf/JFreeChart to graph data
- Ability to fetch Inca data in HTML or XML format via REST URLs **new**
- Allow “run nows” from the Inca web status pages **new**



New features

Depot peering provides fault tolerance



“Approval” mode allows system administrators greater control over testing on their resources

```
Proposed changes to Inca deployment

Suite sampleSuite
[X] update openssh_version
[X] add posixPs_test

space: toggle selection  ^Y: show details  ^S: save/quit  ^Q: quit
arrow-up, k: move up    arrow-down, j: move down
arrow-left: previous suite  arrow-right: next suite
```

Software status and deployments

Current software version: 2.5
(available from Inca website)

<http://inca.sdsc.edu>



TeraGrid[™]



UC Grid

SDSC SAN DIEGO SUPERCOMPUTER CENTER



Inca TeraGrid deployment

- Running since 2003
- Resource registration in information services (IIS)
- Total of ~2200 tests running on 18 login nodes, 2 grid nodes, and 3 servers
- Email notifications for critical services
- Cross-site tests
- CA certificate and CRL checking



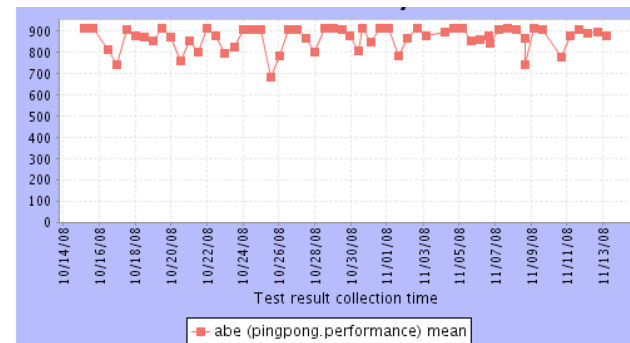
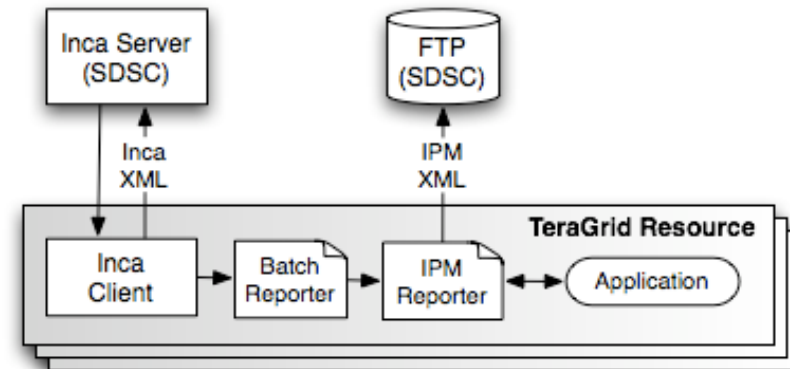
Screenshot of Inca status pages for TeraGrid

<http://inca.teragrid.org/>

*IIS paper to be given on Friday at GCE'09
by JP Navarro*

Using Inca and IPM to measure performance variation on TeraGrid

- Joint work with Nick Wright (LBL) and PMaC group
- Performance variation makes code optimization challenging
- Benchmarks: Paratec, WRF, PingPong, and HPCC (naturally ordered ring and random ring)
- Machines: NCSA Abe, NICS Kraken, and TACC Ranger
- 256 processors twice a day during March 2009
- Variance due to network contention; higher for WRF



Mean MPI ping pong bandwidth history

Wright, N., Smallen, S., Olschanowsky, C., Hayes, J., Snaveley, A. "Measuring and Understanding Variation in Benchmark Performance", *published in HPCMP UGC 2009*

Inca UC Grid deployment

- University of California campuses
- Lead by UCLA
- First deployed in August and now includes 4 campuses
- Total of 71 tests running on 5 portal servers, 4 appliance nodes, and the myproxy server



<http://inca.ucgrid.org>

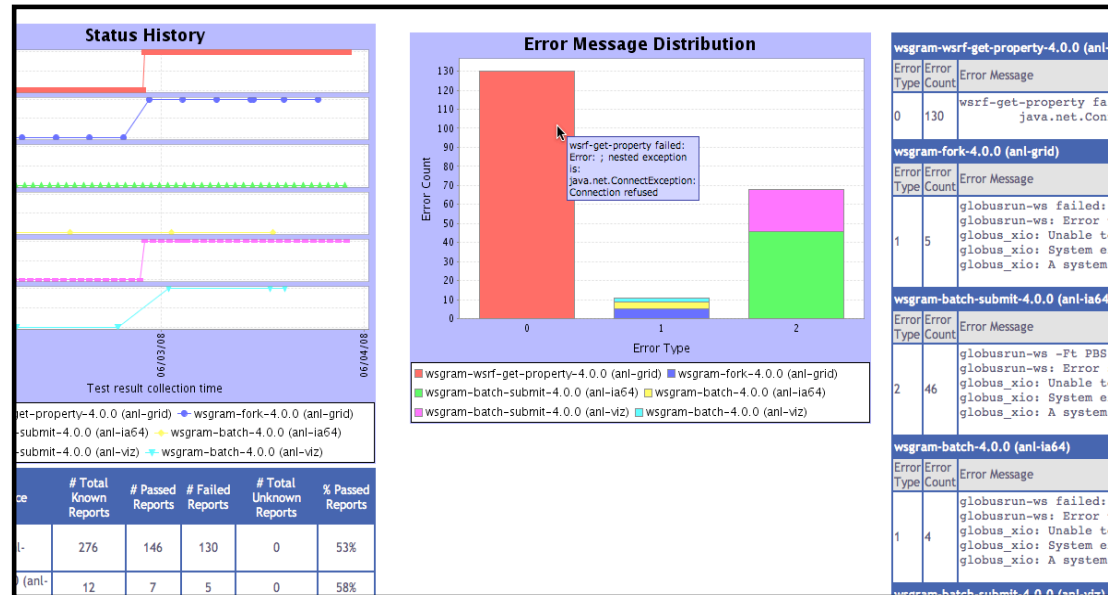
Inca monitoring benefits end users



“Inca reported errors mirror failures we’ve observed and as they are addressed we’ve noticed an improvement in TeraGrid’s stability.”


-- Suresh Marru (LEAD developer)

- Tests resources and services used by LEAD. E.g.
 - Pings service every 3 mins
 - Verifies batch job submission every hour
- Automatically notifies admins of failures
- Show week of history in custom status pages



Benefits of using Inca

- Detect problems before the users notice them
- Easy to write and share tests and benchmarks
- Easy to deploy and maintain
- Flexible and comprehensive displays



Inca Information

- Announcements:
inca-users@sdsc.edu
 - Email:
inca@sdsc.edu
 - Website:
<http://inca.sdsc.edu>
- Supported by:

