User-level Grid Functionality Testing with Inca v2.0

Jim Hayes jhayes@sdsc.edu

SC|06 Tampa, Florida, November 12-17, 2006

Is the User's Grid Working?

- Can user X run application[s] Y on Grid[s] Z? Access dataset[s] N?
  - Can user login?
  - Are Grid services the application[s] use available? Compatible versions?
  - Are dataset[s] N accessible to user X? Credentials?
  - …

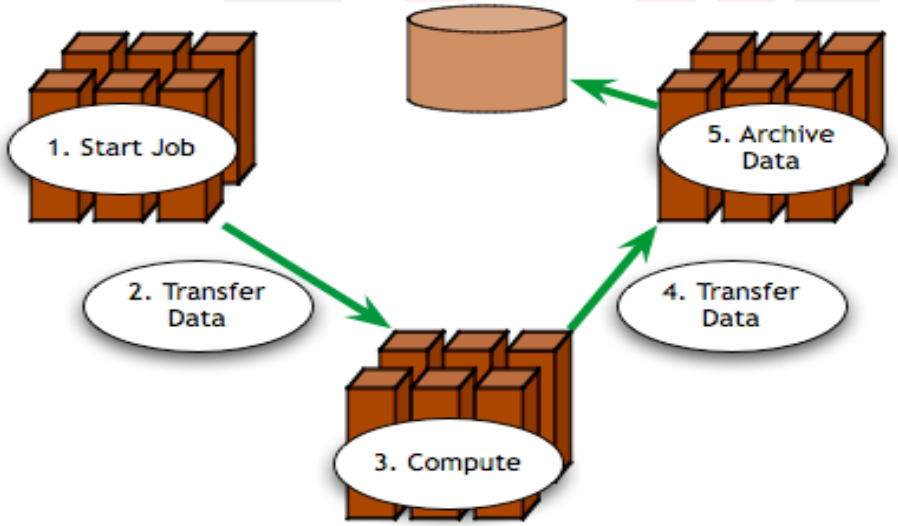

Testing a Grid

1. Iteratively define a set of concrete requirements
2. Write tests to verify requirements
3. Periodically run tests and collect data
4. Publish data

Inca aims to automate steps 3 and 4

What type of testing?

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

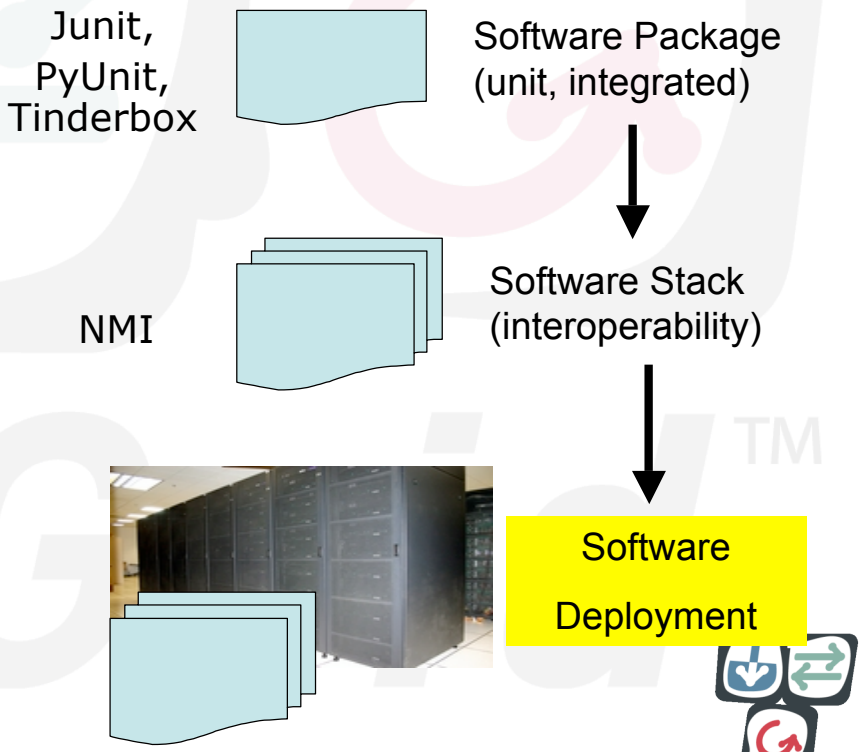

Who are the consumers?

- Grid/VO management
  - Responsible for designing & maintaining requirements
  - Verify requirements are fulfilled by resource providers
- System administrators
  - Notified of problems
  - Enough information to understand context of problem
- End users
  - View results and compare to problems they are having
  - Debug user account/environment issues
  - Feedback to Grid/VO

Inca

Inca is a framework for the automated testing, benchmarking and monitoring of Grid resources

Inca:

- schedules execution of information gathering programs (reporters)
- collects, archives, publishes data

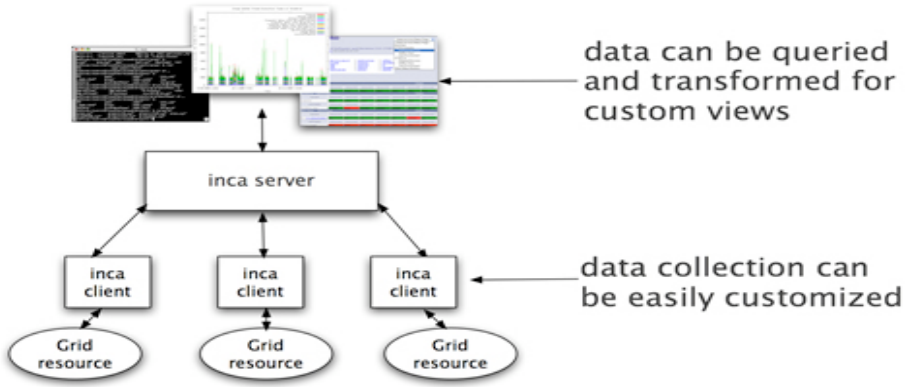

Related Grid monitoring tools

BIG BROTHER™

CLUMON
The Cluster Monitoring System

Ganglia
.sourceforge.net

GridICE
the eyes of the Grid

Hawkeye

the globus toolkit
MDS

MonALISA
MONitoring Agents using a Large Integrated Services Architecture

Nagios®

Inca's primary objective: user-level Grid functionality testing and performance measurement

SC06
POWERFUL BEYOND IMAGINATION

SC|06 Tampa, Florida, November 12-17, 2006

TeraGrid™

Unique features of Inca

- Debugging a deployment
  - Runs under a regular user account
  - Flexibly expresses results
  - Captures test context
  - Securely re-runs tests
  - Archives full reports
  - Tests can be run outside Inca framework

Unique features of Inca (cont.)

- Compares results to a specification
- Easily and securely configured
  - Data collection
  - Installation
- Profiles and logs Inca component resource use

Outline

- Inca in use
- Architecture overview
- Project status

Inca In Use

- Inca Version 1 first put into production in 2004
- Inca Version 2 currently beta
- *Both versions of Inca used in production environments*

GEON

DEISA

NGS

CINECA

TeraGrid™

Inca in use: TeraGrid software stack V&V

- TeraGrid - an “enabling cyberinfrastructure” for scientific research
  - ANL, 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

TeraGrid™

Inca 2 TeraGrid Deployment: CTSSv3

32 packages:
Globus, SRB, Condor-G, MPICH, Softenv, etc.

123 Inca tests:
45 unit/functionality tests
- **Services**: BLAS, uberftp, etc.

29 compatible version tests
- **Version**: HDF, Condor-g, etc.

49 cross-site functionality tests
- **Cross-site**: Globus GRAM, GridFTP, gsissh

Software Stack Summary View

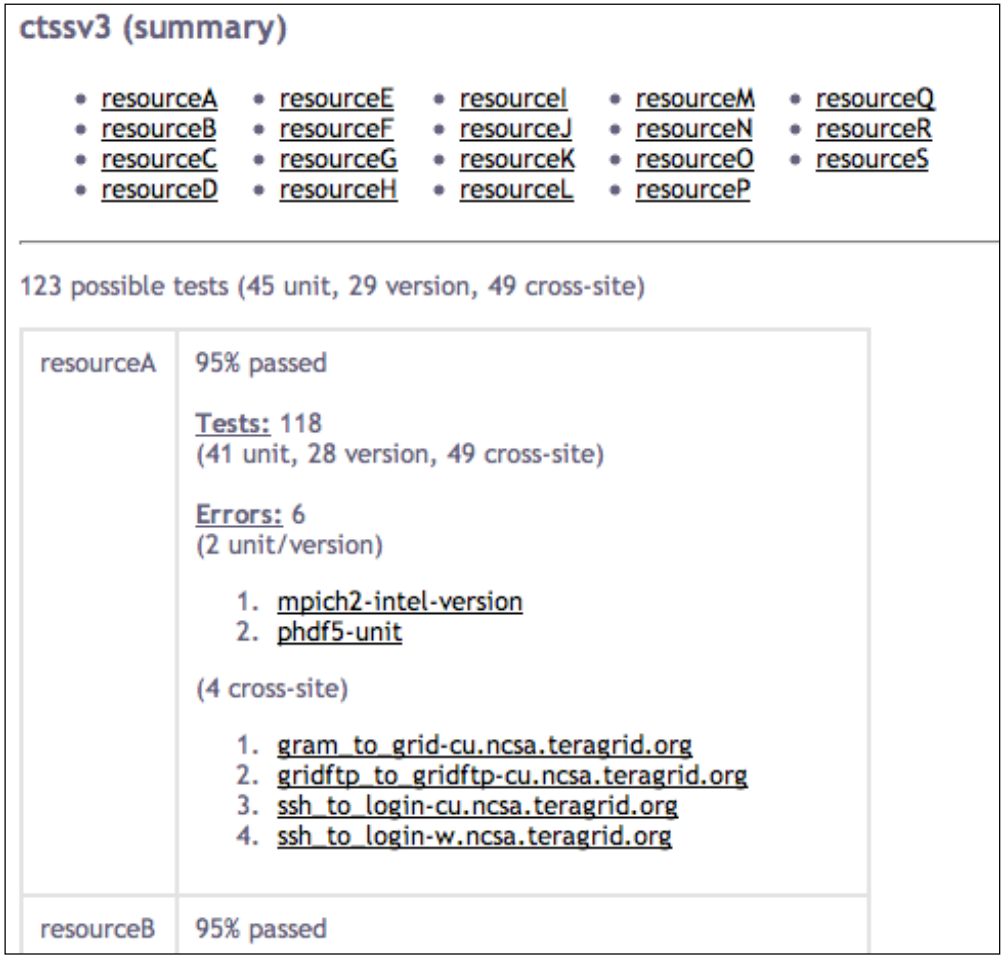

ctssv3 (summary)

- resourceA
- resourceB
- resourceC
- resourceD
- resourceE
- resourceF
- resourceG
- resourceH
- resourceI
- resourceJ
- resourceK
- resourceL
- resourceM
- resourceN
- resourceO
- resourceP
- resourceQ
- resourceR
- resourceS

123 possible tests (45 unit, 29 version, 49 cross-site)

resourceA	95% passed Tests: 118 (41 unit, 28 version, 49 cross-site) Errors: 6 (2 unit/version) 1. mpich2-intel-version 2. phdf5-unit (4 cross-site) 1. gram_to_grid-cu.ncsa.teragrid.org 2. gridftp_to_gridftp-cu.ncsa.teragrid.org 3. ssh_to_login-cu.ncsa.teragrid.org 4. ssh_to_login-w.ncsa.teragrid.org
resourceB	95% passed

Software Stack Detail View

Resources →

SW packages ↓

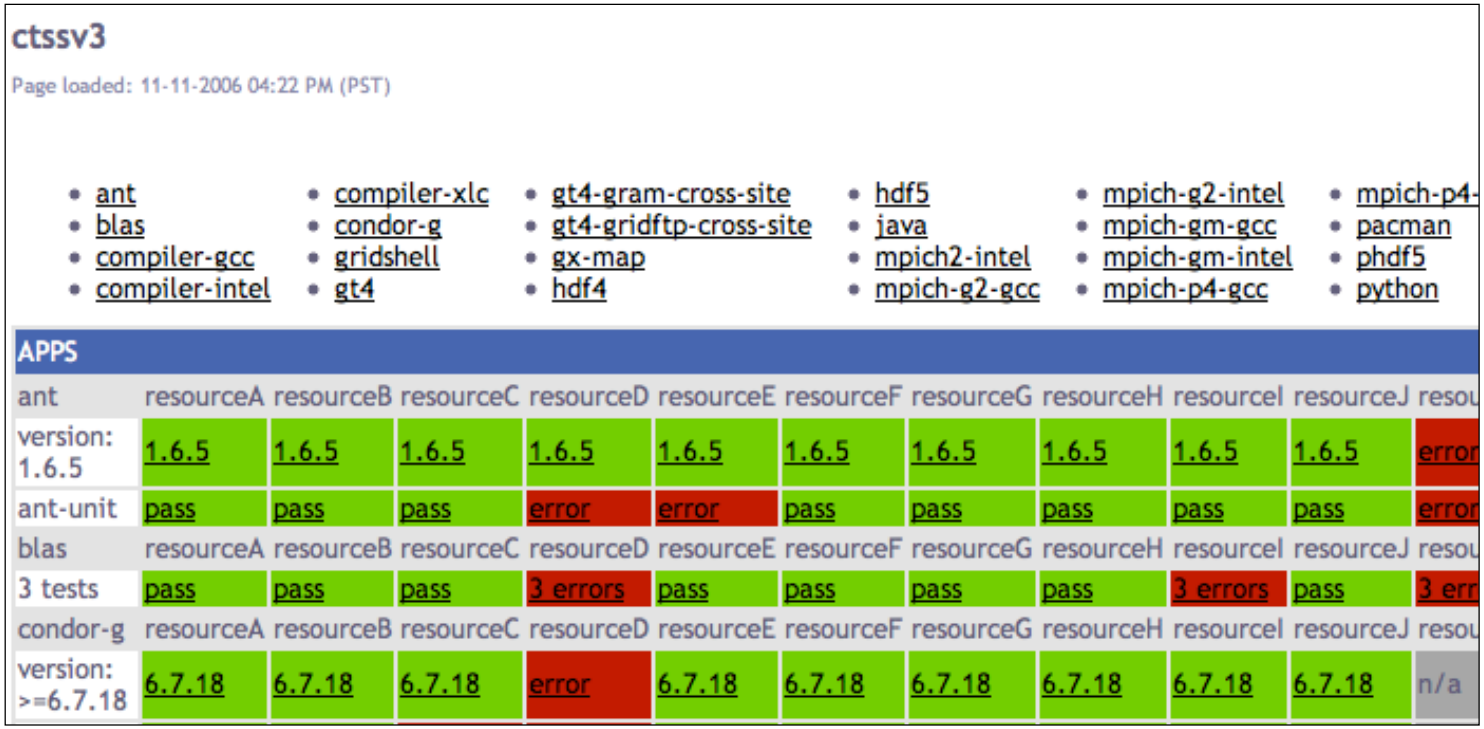

ctssv3

Page loaded: 11-11-2006 04:22 PM (PST)

- ant
- blas
- compiler-gcc
- compiler-intel
- compiler-xlc
- condor-g
- gridshell
- gt4
- gt4-gram-cross-site
- gt4-gridftp-cross-site
- gx-map
- hdf4
- hdf5
- java
- mpich2-intel
- mpich-g2-gcc
- mpich-g2-intel
- mpich-gm-gcc
- mpich-gm-intel
- mpich-p4-gcc
- mpich-p4-
- pacman
- phdf5
- python

APPS

ant	resourceA	resourceB	resourceC	resourceD	resourceE	resourceF	resourceG	resourceH	resourceI	resourceJ	resou
version: 1.6.5	1.6.5	1.6.5	1.6.5	1.6.5	1.6.5	1.6.5	1.6.5	1.6.5	1.6.5	1.6.5	error
ant-unit	pass	pass	pass	error	error	pass	pass	pass	pass	pass	error
blas	resourceA	resourceB	resourceC	resourceD	resourceE	resourceF	resourceG	resourceH	resourceI	resourceJ	resou
3 tests	pass	pass	pass	3 errors	pass	pass	pass	pass	3 errors	pass	3 err
condor-g	resourceA	resourceB	resourceC	resourceD	resourceE	resourceF	resourceG	resourceH	resourceI	resourceJ	resou
version: >=6.7.18	6.7.18	6.7.18	6.7.18	error	6.7.18	6.7.18	6.7.18	6.7.18	6.7.18	6.7.18	n/a

Single Result View

Result:

did not complete

```
globusrun failed:
GRAM Authentication test failure: connecting to the
job manager failed. Possible reasons: job
terminated, invalid job contact, network problems,
...
```

Reporter details:

reporter name	grid.middleware.globus.unit.gatekeeper (click name for more info)

Execution information:

ran at	11-10-2006 12:39 PM (PST)
age	22 hours 44 minutes
cron	?=36 ?=14 * * *
ran on (hostname)	resource.teragrid.org
memory usage (MB)	19.1094
cpu time (secs)	3.46289
wall clock time (secs)	183.855

Outline

- Inca in use
- Architecture overview
- Project status

Inca Components

- Inca Control System
  - Reporter Managers manage Inca on a single host
  - Agent installs and oversees Reporter Managers
  - Incat provides an administration GUI
- Inca Data System
  - Reporters perform tests; collected in external repositories
  - Depot stores and retrieves test results (DB)
  - Consumer provides web-based access to test results

1. The user creates a *suite* using **incat** and submits it to the **agent**

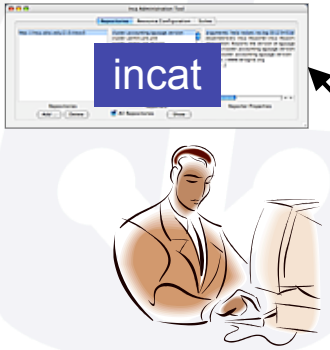

2. The **agent**
   - fetches *reporters* from the **reporter repository**
   - creates a **reporter manager** on each resource
   - sends the suite and reporters to each **reporter manager**.

4. **Data consumers** display collected data (*reports*) by querying the **depot**

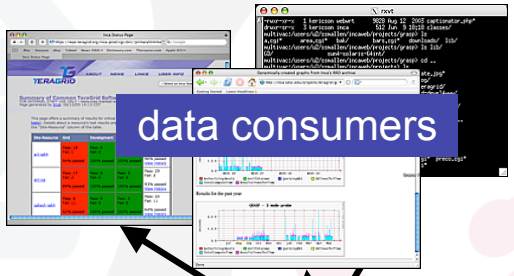

3. Each **reporter manager** executes *reporters* according to its schedule and sends results (*reports*) to the **depot**

TeraGrid™

Inca Component Design Goals

- Minimal impact on monitored resources
  - Most administrative duties handled by Agent
  - Component profiling; Reporter profiling/limits
- Easy installation and maintenance
  - Centralized administration control
  - Automated staging of Reporter Managers and Reporters
- Flexible reporter scheduling and configuration
  - On-demand and periodic scheduling
  - Reporters independent of Inca deployment

Inca Component Design Goals (cont.)

- Security
  - Certificate-based authentication by all components
  - Provide credentials to reporters that need it
- Flexible data content
  - No required schema for test output
- Efficient storage and retrieval of data
  - Storage in a selection of databases via Hibernate
  - Schema designed to reduce redundant data
- Flexible access to data
  - Access to all data via SQL queries
  - Depot protocol provides predefined queries; extensible
  - Depot retains history of reporter output
  - Queries via Inca protocol and web service interface

Inca Data Consumer

ctssv3 (summary)

- resourceA
- resourceB
- resourceC
- resourceD
- resourceE
- resourceF
- resourceG
- resourceH
- resourceI
- resourceJ
- resourceK
- resourceL
- resourceM
- resourceN
- resourceO
- resourceP
- resourceQ
- resourceR
- resourceS

123 possible tests (45 unit, 29 version, 49 cross-site)

resourceA	95% passed Tests: 118 (41 unit, 28 version, 49 cross-site) Errors: 6 (2 unit/version) 1. mpich2-intel-version 2. phdf5-unit (4 cross-site) 1. gram_to_grid-cu.ncsa.teragrid.org 2. gridftp_to_gridftp-cu.ncsa.teragrid.org 3. ssh_to_login-cu.ncsa.teragrid.org 4. ssh_to_login-w.ncsa.teragrid.org
resourceB	95% passed

Inca Reporter

- Executable program that measures some aspect of the system or installed software
- Requirements:
  - Supports specific command-line options
  - Writes XML (Inca Reporter schema) to stdout
- Extensive Library support for perl scripts

Example: openssh version

```
use Inca::Reporter::Version;
my $reporter = new Inca::Reporter::Version(
  version => 1.8,
  description => 'Reports the version of openssh',
  url => 'http://www.openssh.org',
  package_name => 'openssh'
);
$reporter->processArgv(@ARGV);
$reporter->setVersionByExecutable('ssh -V', 'OpenSSH_([\w\.]+)|GSI ([\w\.]+)');
$reporter->print();
```

Incat Administration Tool

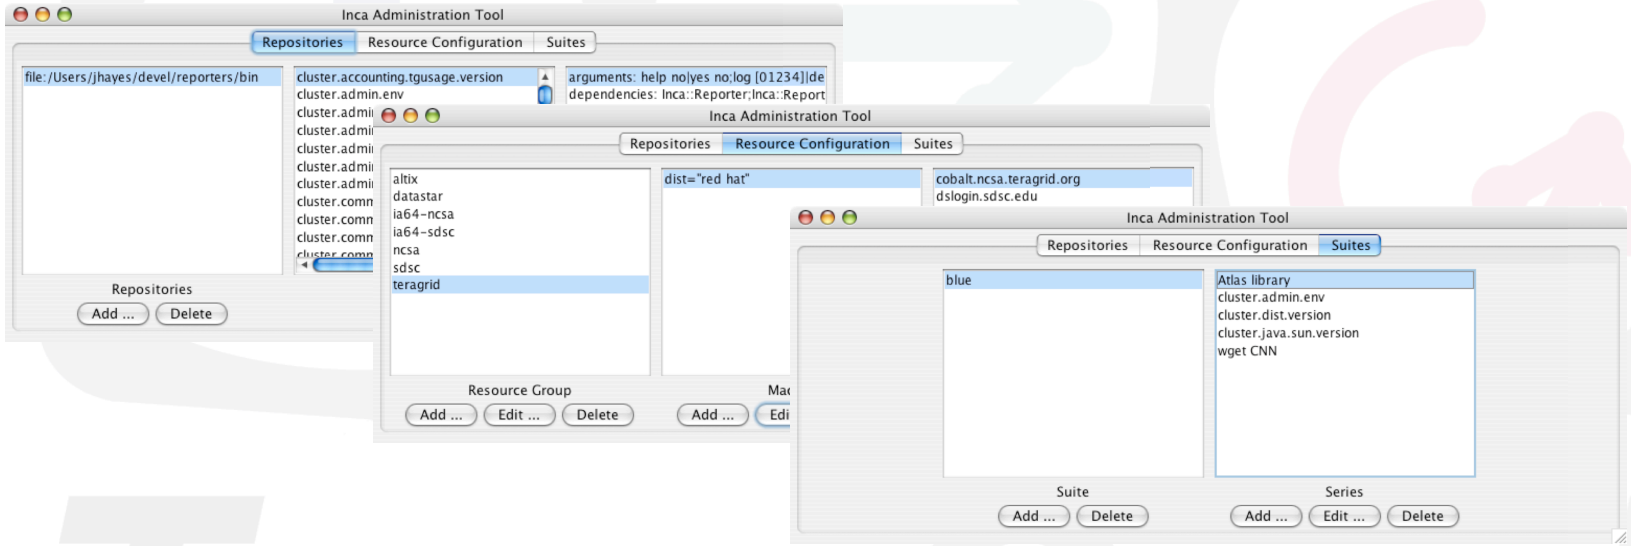

- Centralized configuration of Inca installation
- Reporter repositories, resources/hosts, suites

Outline

- Inca in use
- Architecture overview
- Project status

Inca v2 Development

- Version 1 being phased out
- Version 2 beta released 11/06
- Production version available “soon”
- Future work
  - Scalable design
  - Improved displays
  - Extended package support

Scalable design

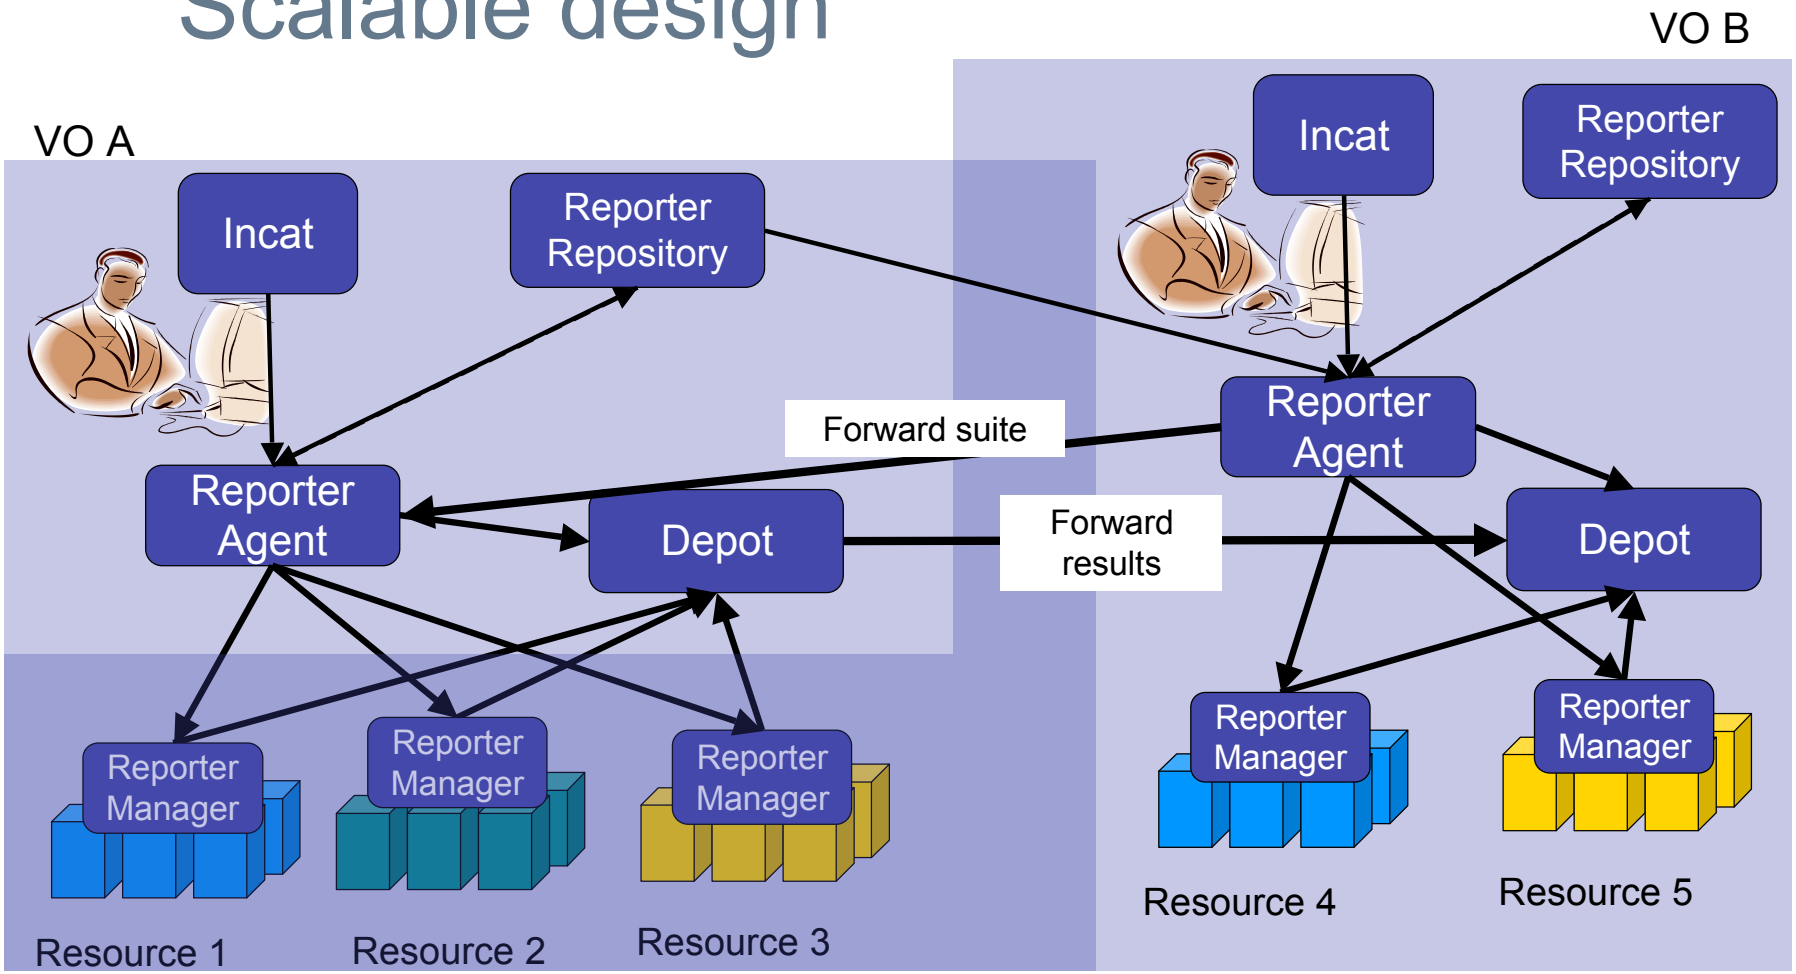

Historical Graphs

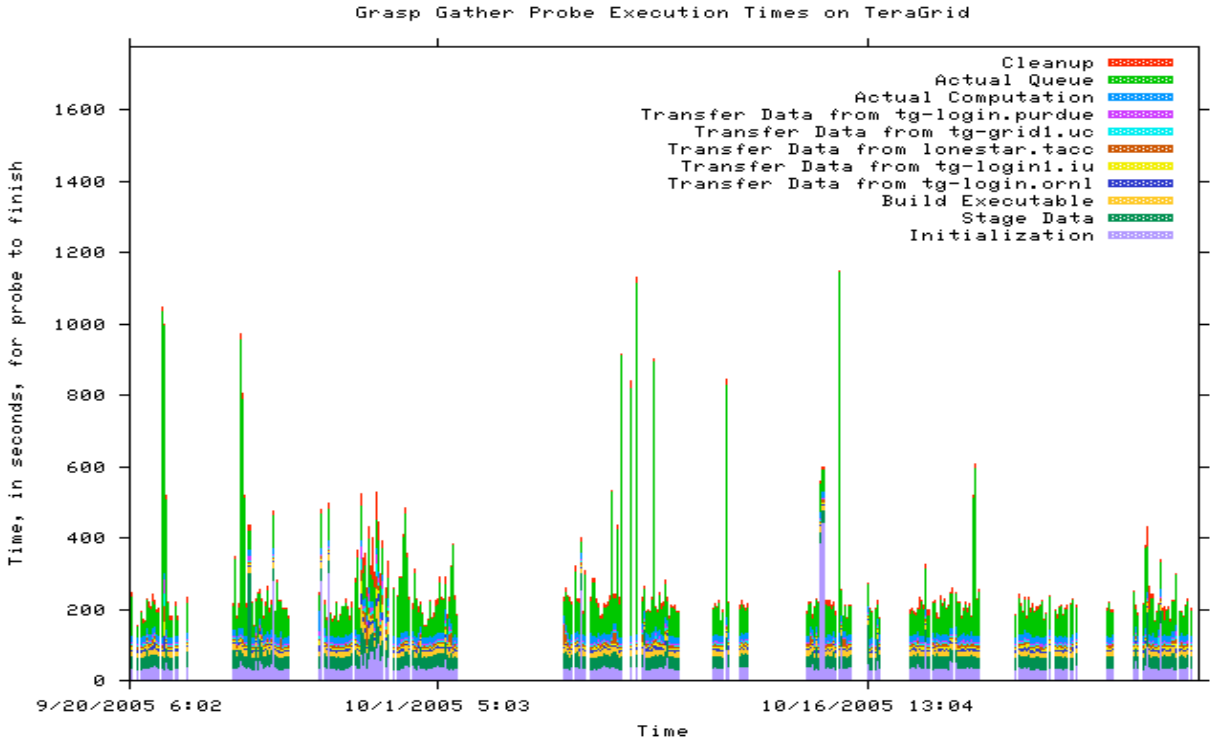

Historical Graphs

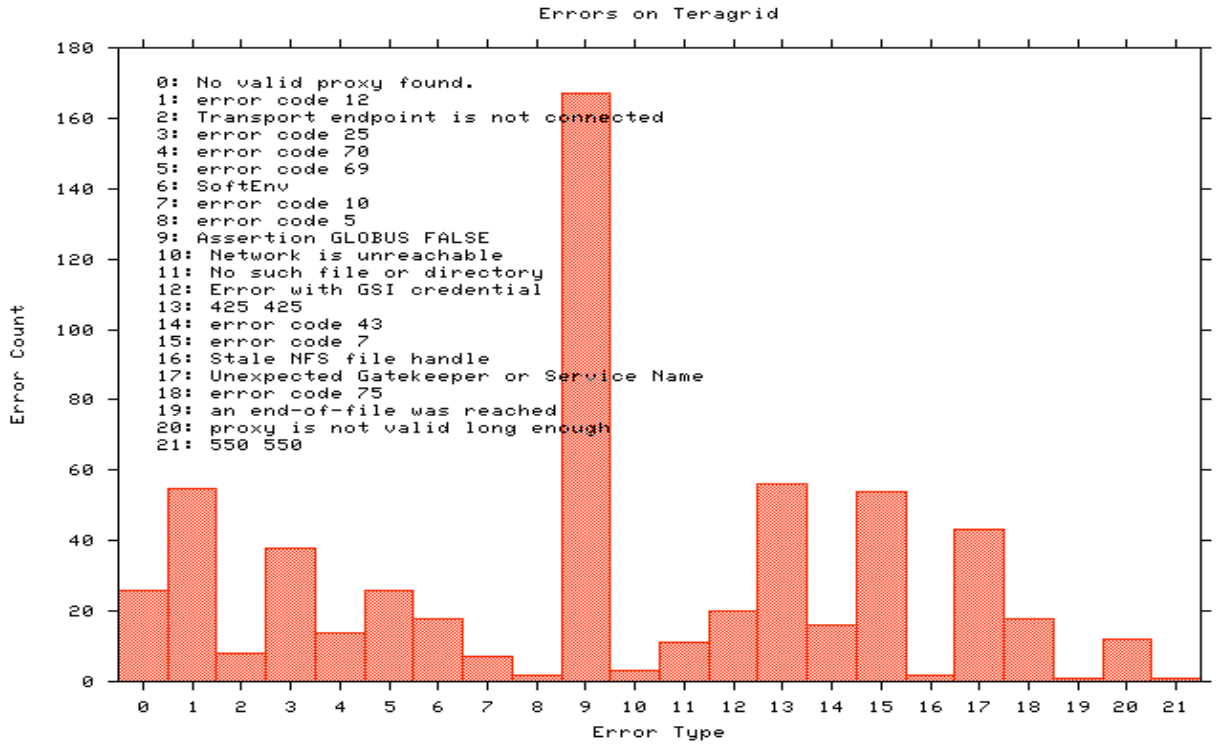

Error tracking over time

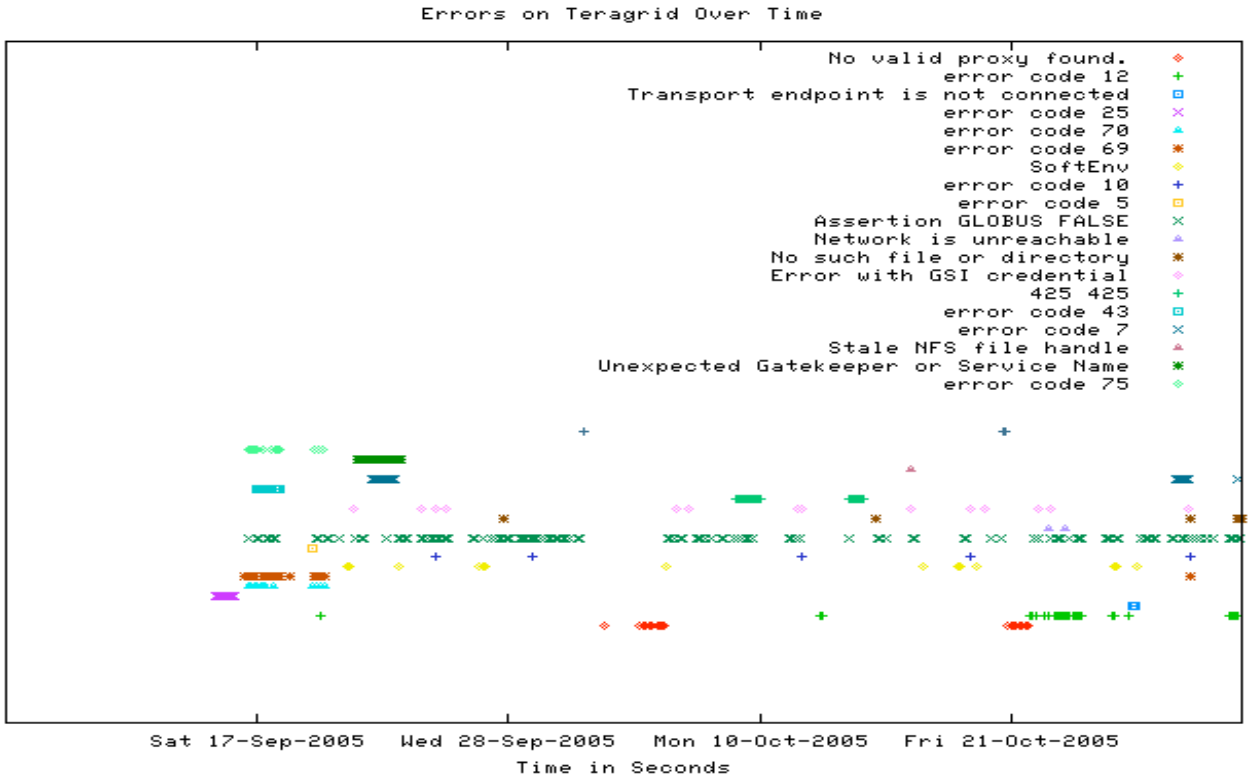

Extended Package Support

- Reporter Repositories can presently contain scripts, perl modules, and .tar.gz packages
- Add support for .rpm, other formats
- Support retrieval from CPAN

More information

- Inca Web Page:
  http://inca.sdsc.edu
- Email
  - inca@sdsc.edu