

User-level Grid monitoring with Inca 2

Shava Smallen
ssmallen@sdsc.edu

June 25, 2007



HPDC Grid Monitoring Workshop
June 25, 2007



TeraGrid

- Origins: national supercomputer centers, funded by the NSF
- 9 TeraGrid sites, 18 resources
- Mix of Architectures:
 - ia64, ia32: LINUX
 - Cray XT3
 - Alpha: True 64
 - SGI SMPs



- Connected via dedicated multi-Gbps links
- 1000s of CPUs, > 250 teraflops
- > 30 petabytes of online and archival data storage
- Coordinated user environment across heterogeneous resources
 - CTSS (Coordinated TeraGrid Software & Services)



HPDC Grid Monitoring Workshop
June 25, 2007



User-level Grid monitoring

- Testing and performance measurement from a generic, impartial user's perspective in order to detect and fix Grid infrastructure problems before the user's notice them.
- User-level Grid monitoring system:
 - Runs from a standard user account
 - Executes using a standard GSI credential
 - Uses tests that are developed and configured based on user documentation
 - Verifies user-accessible Grid access points
 - Centrally manages monitoring configuration
 - Automates periodic execution of tests
 - Easily updates and maintains monitoring deployment



HPDC Grid Monitoring Workshop
June 25, 2007

SDSC



Inca

- Provides user-level monitoring of Grid functionality and performance
- Features:
 - Collects wide variety of monitoring results
 - Captures context of monitoring result as it executes
 - Eases the writing and deploying of new tests or benchmarks
 - Supports sharing of tests and benchmarks
 - Stores and archives monitoring results
 - Securely manages short-term proxies
 - Measures system impact of tests and benchmarks

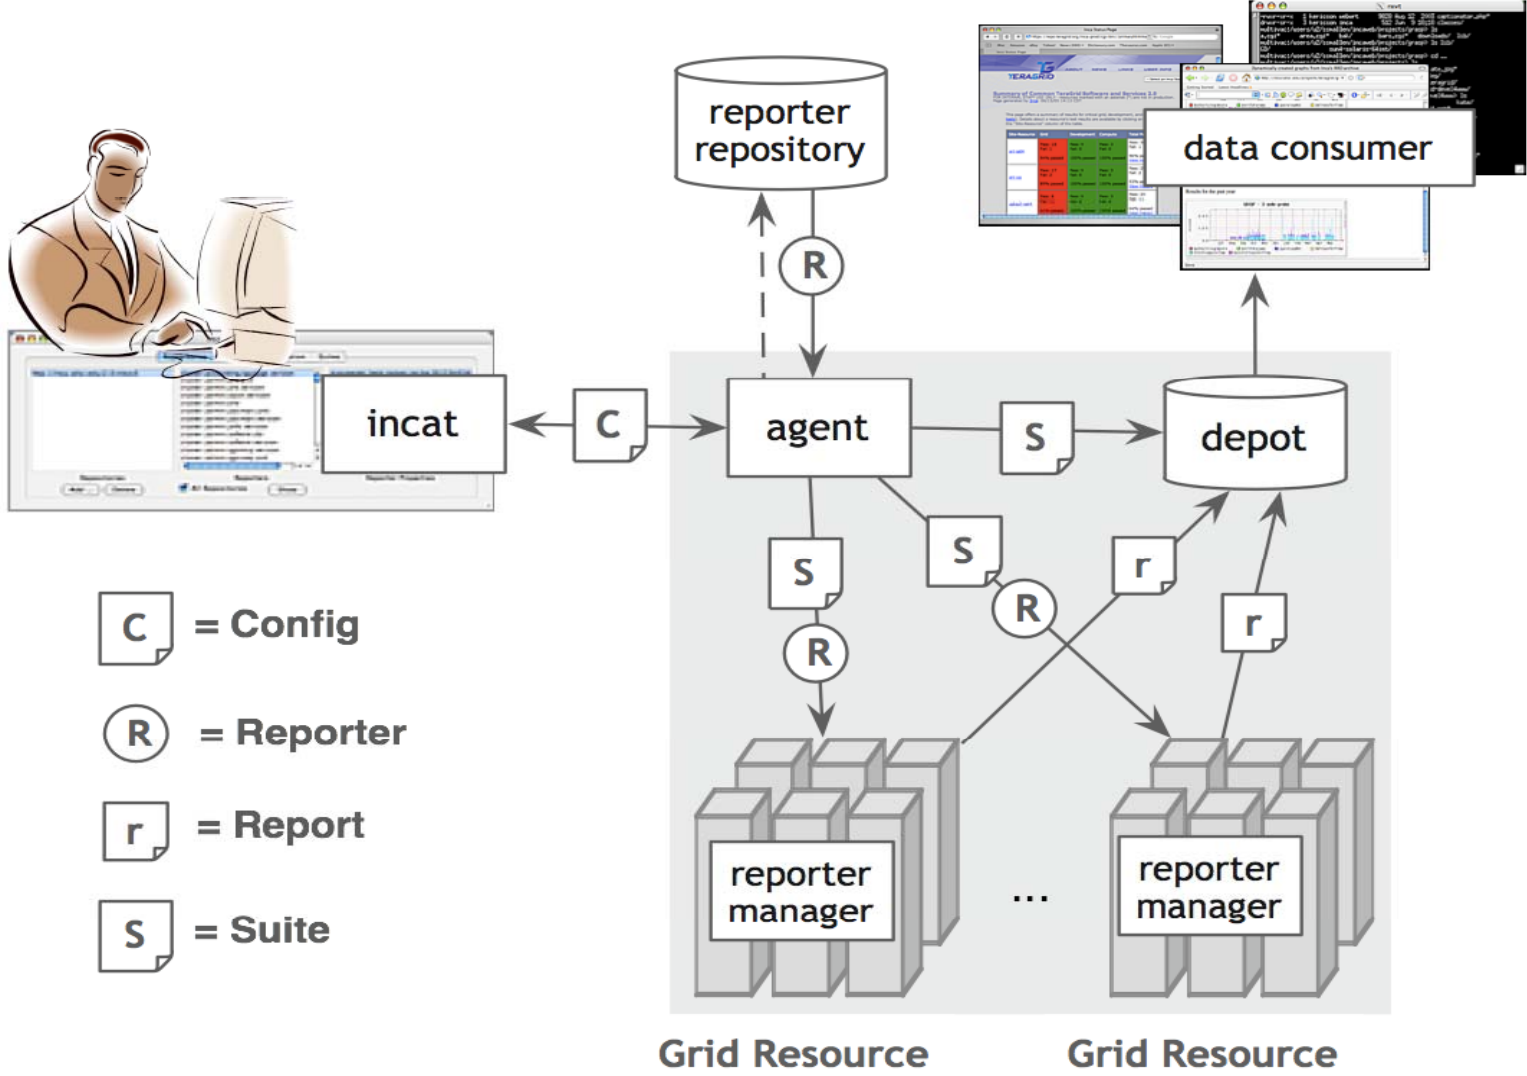


HPDC Grid Monitoring Workshop
June 25, 2007

SDSC



Inca Architecture



HPDC Grid Monitoring Workshop
June 25, 2007

SDSC



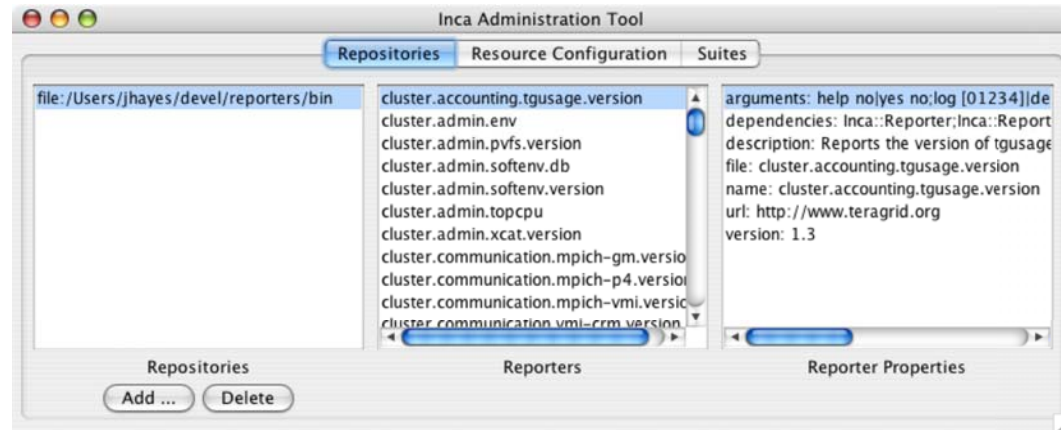
Collecting Monitoring Data

- Reporters
 - 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
 - Supports multiple types of data
 - Extensive library support for perl scripts
 - Most reporters < 30 lines of code
 - Independent of other Inca components



Sharing Reporters

- Repositories: collection of reporters available via a URL
 - Supports package dependencies (Perl modules, Makefile, autoconf)
 - Packages versioned to allow for automatic updates
 - Inca repository contains 157 reporters
 - Version, unit test, performance benchmark reporters
 - Grid middleware and tools, compilers, math libraries, data tools, and viz tool



Screenshot of a repository using Inca GUI tool



Centralized configuration and deployment

- Incat
 - GUI interface to enable a large number of monitoring results to be collected with a minimum of effort
 - Configure the reporters to execute on a set of resources
 - Configuration stored in a XML file and sent to Agent
- Agent
 - Implements the configuration specified by Inca administrator
 - Stages and launches a reporter manager on each resource
 - Sends package and configuration updates



Storing data

- Depot
 - Stores configuration information and monitoring results
 - Uses relational database backend via Hibernate
 - Provides full archiving of reporter output
 - Supports SQL queries and provides predefined queries for latest monitoring results, report instance, and report history
 - Supports notifications



Displaying and publishing data

- Data Consumer
 - Web application that queries and displays monitoring data
 - Packaged with Jetty
 - JSP tags to query data and format using XSL
- Web services
 - Query data from depot and return as XML



Inca in Use: TeraGrid

- Currently monitoring all 18 allocated TeraGrid resources
 - Monitoring of CTSSv3
 - Monitoring of CTSSv4 (in progress)
 - Grid jobs (Globus gatekeeper logs)
 - CA certificate and CRL checking (notify if 2 weeks from expiration)
 - Resource registration in MDS

ctssv3
Page loaded: 05-29-2007 06:30 PM (PDT)

n/a	does not apply to resource		test
	missing (not yet executed)	pkgWait	wai
pass	passed	incaWait	wai
error	error	incaErr	inca

- ant
- blas
- compiler-gcc
- compiler-intel
- compiler-xlc
- condor-g
- gridshell
- gsissh
- gt4
- gt4-gr
- gt4-gr

APPS

ant	anl-ia64
version: 1.6.5	1.6.5
ant-unit	pass
blas	anl-ia64
blas-level1	pass
blas-level2	pass
blas-level3	pass
condor-g	anl-ia64
version: >=6.7.18	6.7.18
condorg-condorq	pass

Details for "all2all:gram_to_gatekeeper.bigred.iu.teragrid.org" series

Result:
completed

Reporter details:

reporter name	grid.middleware.globus.unit.gatekeeper (click name for more info)
reporter version	3

Execution information:

ran at	06-23-2007 02:57 AM (PDT)
age	10 hours 46 minutes
cron	?=57 ?=4 ***
ran on (hostname)	tg-login1.uc.teragrid.org
memory usage (MB)	19.1406
cpu time (secs)	0.800782
wall clock time (secs)	1.16846

Input parameters:

help	no
host	gatekeeper.bigred.iu.teragrid.org
log	3
verbose	1
version	no

Command used to execute the reporter:

```
% bash -l -c 'set -a; cd /home/inca/inca2install-ia64; cp ~/.soft.v
&& soft-misc ~/.soft.v3.$$ && source ~/.soft.v3.$$cache.sh && expo
PERL5LIB=/home/inca/inca2install-ia64/var/reporter-
packages/lib/perl:${HOME}/inca/install-ia64/lib/perl &&/home/inca/
ia64/var/reporter-packages/bin/grid.middleware.globus.unit.gatekee
host=gatekeeper.bigred.iu.teragrid.org --log=3 --verbose=1 --
-f ~/.soft.v3.$$*';
```

System commands executed by the reporter:

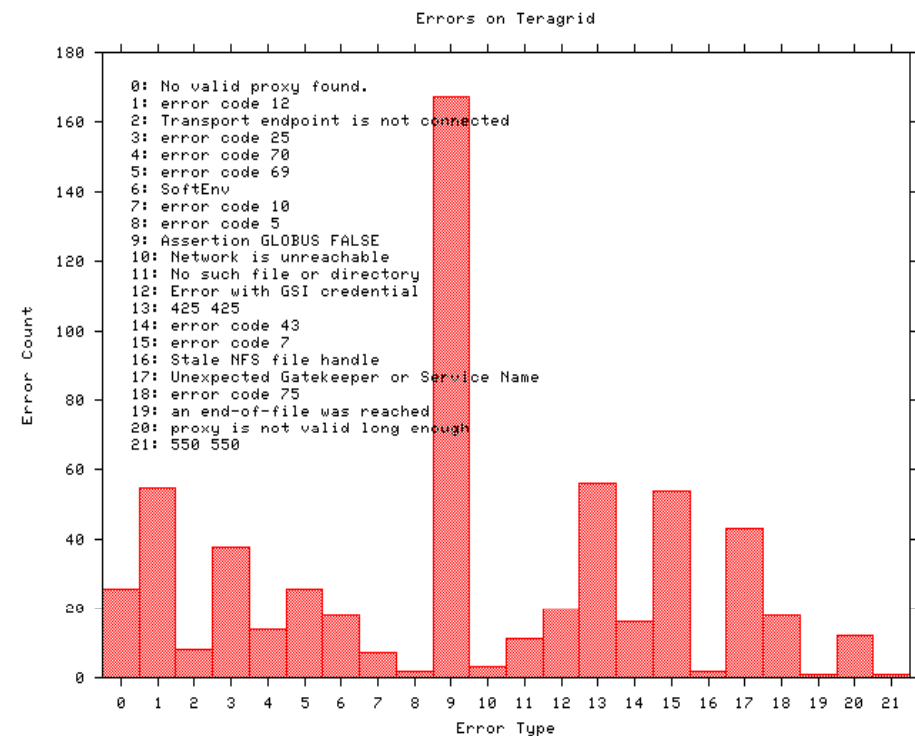
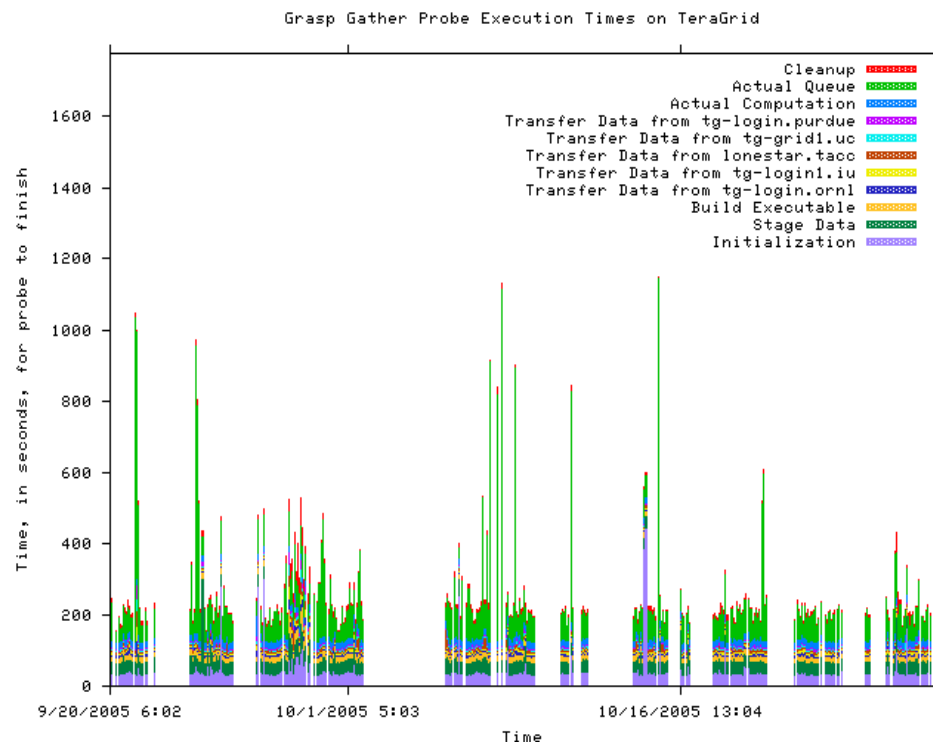
```
% globusrun -a -r gatekeeper.bigred.iu.teragrid.org
```

Note that the reporter may execute other actions in between system commands (e.g., change dire



Inca in use: Grid Assessment Probes

- Set of probes designed to emulate Grid applications
- Deployed using Inca to GEON and TeraGrid



HPDC Grid Monitoring Workshop
June 25, 2007

SDSC



Software Status

Current software version: 2.03
(available from Inca website)
<http://inca.sdsc.edu>

Other Inca deployments:



HPDC Grid Monitoring Workshop
June 25, 2007

SDSC



Summary

- User-level Grid monitoring: Testing and performance measurement from an impartial user perspective to detect problems before the users notice them
- Standalone reporter APIs and repositories make it easy to write and share tests and benchmarks (reporters)
- Centralized configuration enables uniform monitoring and makes it easy to deploy Inca monitoring to a set of resources
- Data consumer and web services interface enable publishing and displaying of Inca monitoring data



HPDC Grid Monitoring Workshop
June 25, 2007

SDSC



More Information

Website:

<http://inca.sdsc.edu>

Announcements:

inca-users@sdsc.edu

Email:

inca@sdsc.edu

Supported by:

SDSC

neon
NATIONAL ECOLOGICAL OBSERVATORY NETWORK



TeraGrid™



HPDC Grid Monitoring Workshop
June 25, 2007

SDSC



TeraGrid™

Sample Reporter

```
use Inca::Reporter::SimpleUnit;
my $reporter = new Inca::Reporter::SimpleUnit(
    name => 'grid.globus.gramPing',
    version => 2,
    description => 'Checks gatekeeper is accessible from local machine',
    url => 'http://www.globus.org',
    unit_name => 'gramPing'
);
$reporter->addDependency('Inca::Reporter::GridProxy');
$reporter->addArg('host', 'gatekeeper host');
$reporter->processArgv(@ARGV);
my $host = $reporter->argValue('host');
my $out = $reporter->loggedCommand("globusrun -a -r $host", 30);
if (!$out) {
    $reporter->unitFailure("globusrun failed: $!");
} elsif($out !~ /GRAM Authentication test successful/) {
    $reporter->unitFailure("globusrun failed: $out");
} else {
    $reporter->unitSuccess();
}
$reporter->print();
```



HPDC Grid Monitoring Workshop
June 25, 2007



Scheduling and Execution

- Reporter manager
 - Manages and schedules the execution of reporters on a single resource
 - Executes under regular user account
 - Monitors reporter system usage and enforces limits
 - Sends monitoring result to a depot

