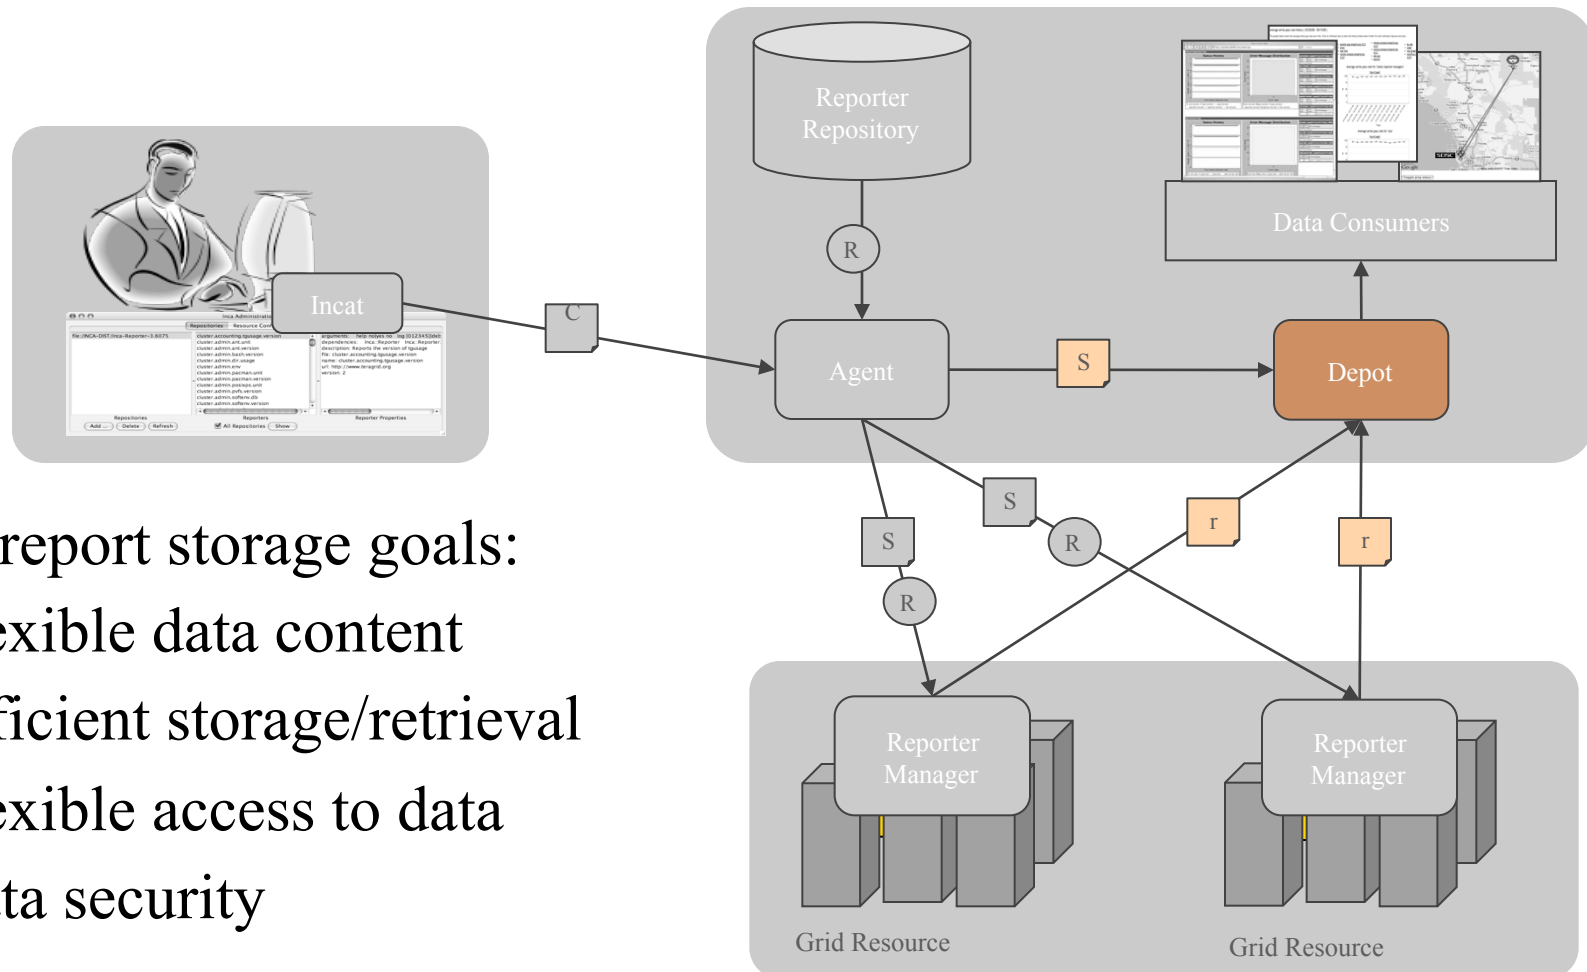

Inside the Inca Depot

Paul Hoover

phoover@sdsc.edu

Inca Workshop
August 26-27, 2010

Depot



Inca report storage goals:

- Flexible data content
- Efficient storage/retrieval
- Flexible access to data
- Data security

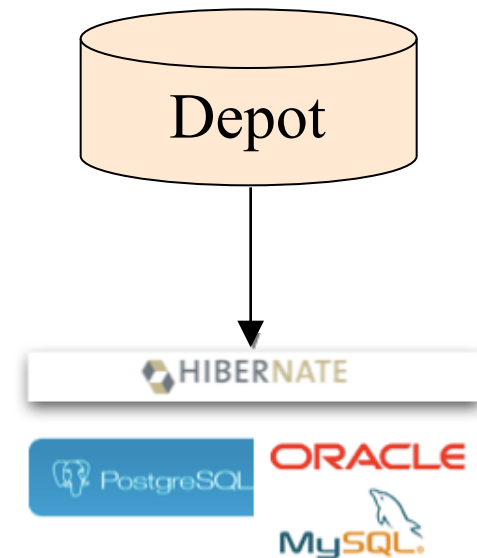
Report Body Provides Flexible Data Content

```
<body>
  <performance ID="hpcc">
    <benchmark ID="8">
      <statistics
        EP-STREAM_Triad_GB_s="2.813"
        G-STREAM_Triad_GB_s="22.505"
        Wall_Mins="74.317"
        EP-DGEMM_GFlop_s="8.936"
        G-HPL_TFlop_s="0.067"
        Random_Ring_Latency_usecs="0.744"
        ProcSpeed_GHz="2.60"
        Cores="8.000"
        HPL_percent="80.639"
        G-Random_Access_GFlop_s="0.057"/>
      </benchmark>
    </performance>
  </body>
```

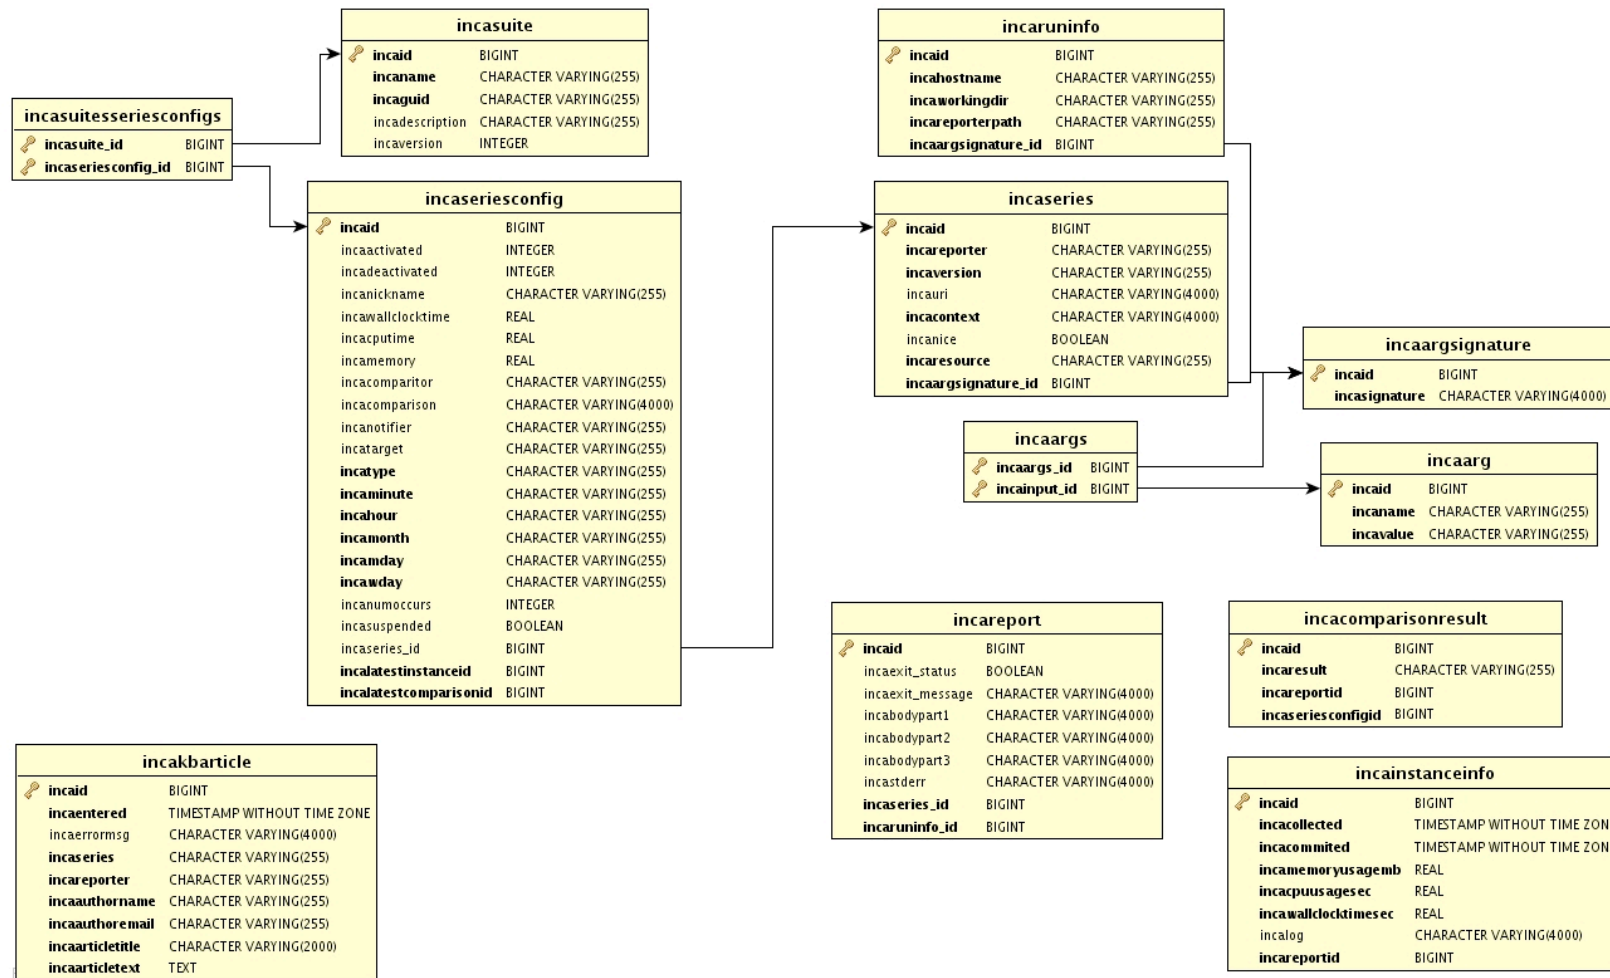
```
<body>
  <package>
    <ID>openssh</ID>
    <version>3.6.1p1</version>
  </package>
</body>
```

Hibernate Supports Efficient Data Access

- Hibernate library provides an abstraction layer to DB applications
- Supports PostgreSQL, MySQL, Oracle, HSQLDB, others
- Provides HQL object-oriented SQL dialect



Database Schema



inca Series Edit Dialog

Incat Series Dialog

Reporter *

Resource *

host @info_service_hosts@

lines 10

log 5

match TeraGrid.*Service

path

port 80

Reset to Defaults

Arguments

Use Version 1 Only Use Latest Version
(Reports whether a service is accessible via http)

Run Every 1 hours From 0 ? * * * *

Nickname @info_service_hosts@:80 (Apache)

Context /bin/bash -l -c'grid.middleware.globus.unit.httpmds -args'

Limits CPU (Secs) Wall (Secs) @maxWallTime@ Memory (MB)

Comparison errorMessage == " || errorMessage =~ 'Inca error'

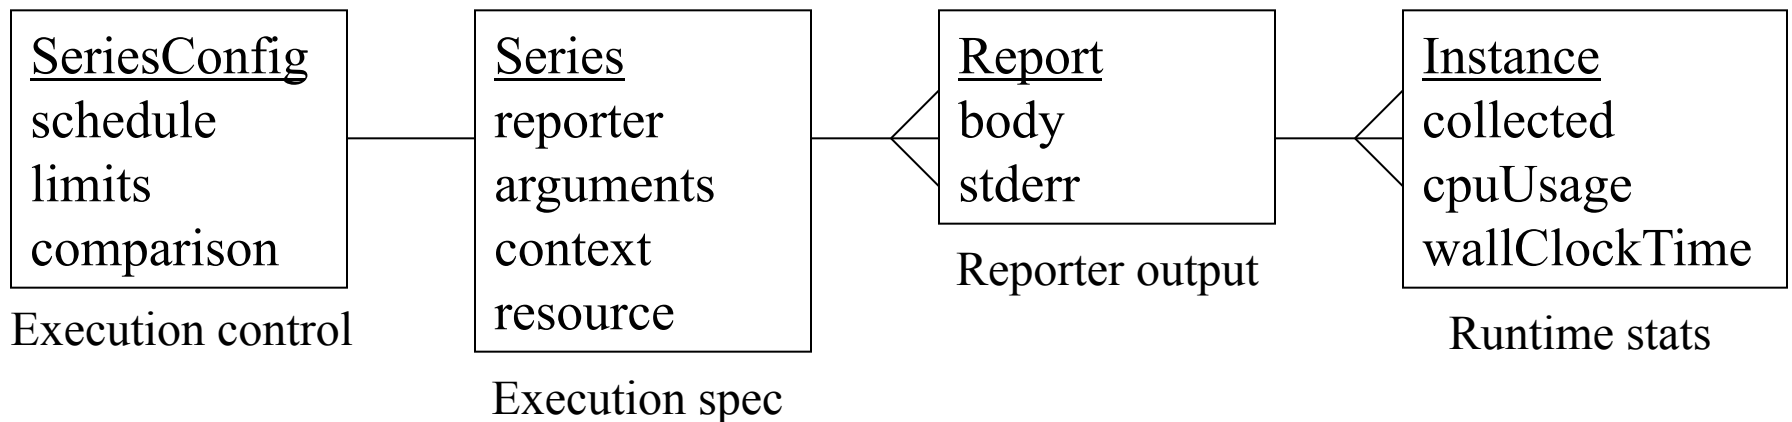
Notification Script EmailNotifier

Script Arguments inca@sdsc.edu

Cancel Ok

Depot Schema Reduces Redundancy

- Report bodies often identical; Depot merges them
- Instance contains info that varies--timestamps, resource usage
- Series defines reporter execution; SeriesConfig manipulates execution and output



Depot Provides Flexible Access to Data

- Depot retains history of reporter output
- Protocol supports direct DB access via HQL queries
- Convenience queries provide easy access to common combinations of data
 - Single instance
 - Latest instance of a set of series
 - Instances collected during a specified time period
 - Success/failure history of a set of series
- Special queries report DB schema, suite list

Depot Customization

- Switch underlying database engine
- Add/modify notifier scripts
- Add report filters

Hibernate Supports Multiple DB Engines

- PostgreSQL, MySQL, Oracle, HSQLDB, others
- Differ in optimizations, tools
- Configured via \$INCA_DIST/etc/hibernate.properties
- Needs appropriate JDBC driver

Settings for a local PostgreSQL database.

hibernate.dialect=org.hibernate.dialect.PostgreSQLDialect

hibernate.query.substitutions=yes 'Y', no 'N'

hibernate.connection.driver_class=org.postgresql.Driver

hibernate.connection.url=jdbc:postgresql:test

hibernate.connection.username=test

hibernate.connection.password=

Settings for a remote MySQL database.

hibernate.dialect=org.hibernate.dialect.MySQLDialect

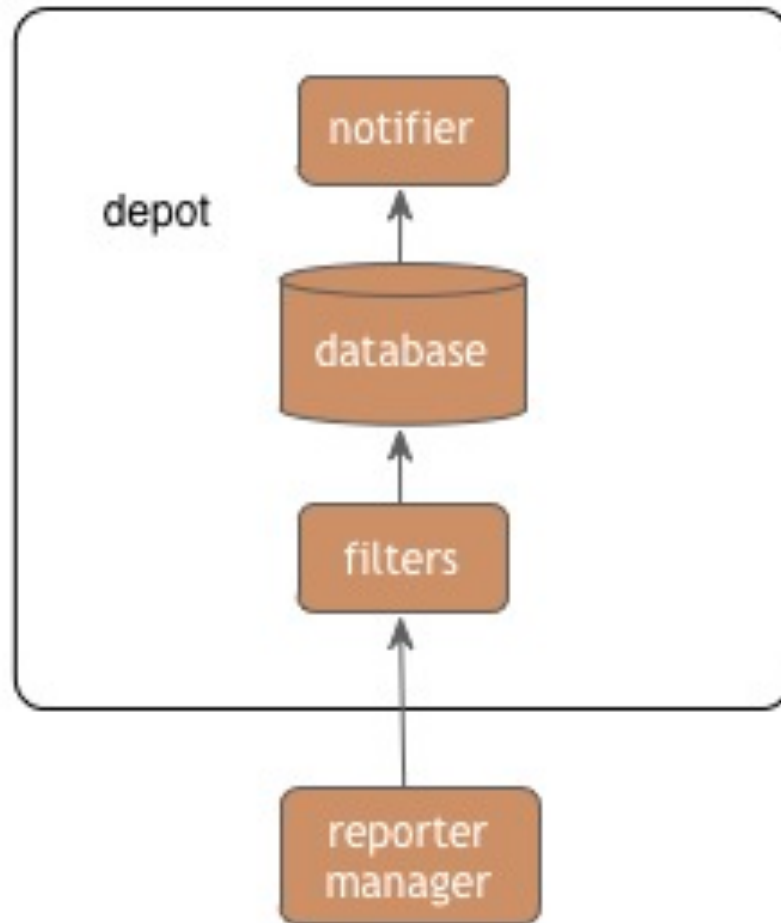
hibernate.connection.driver_class=com.mysql.jdbc.Driver

hibernate.connection.url=jdbc:mysql://192.168.0.6/test

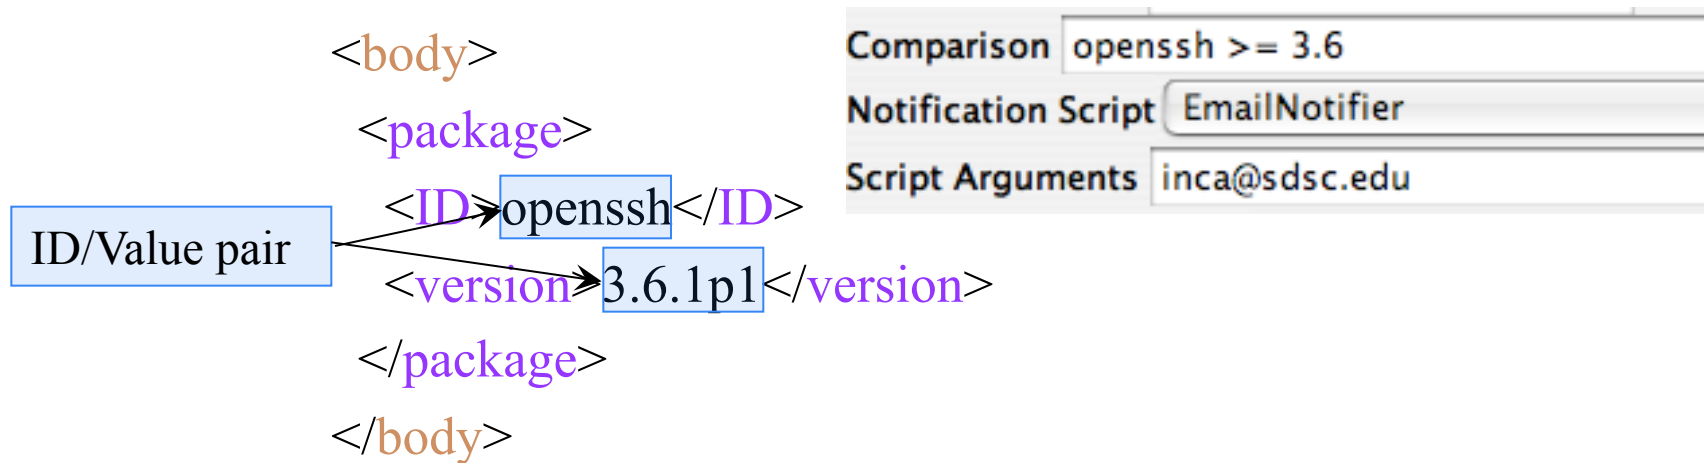
hibernate.connection.username=test

hibernate.connection.password=

Notifiers and Report Filters



Notifiers Report Comparison Success/Failure



- Notifier scripts found in installation `$INCA_DIST/sbin/`
- EmailNotifier, LogNotifier provided; custom notifiers supported
- Invoked on status change (new/changed failure, new success)

LogNotifier

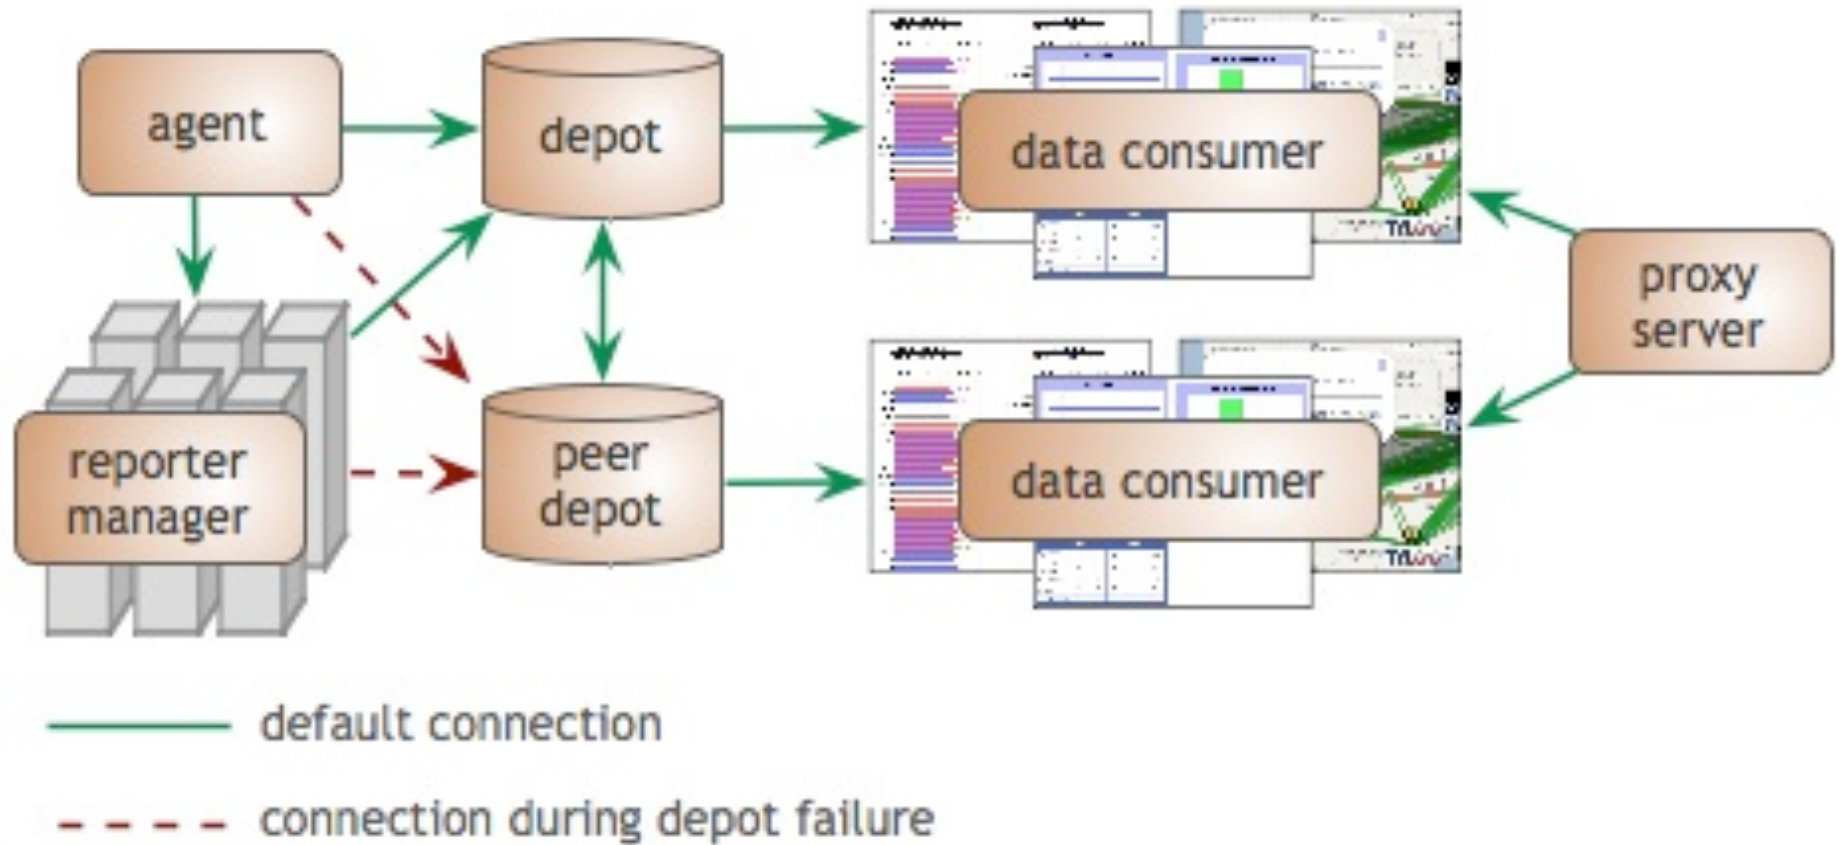
```
use Cwd qw(abs_path getcwd);
my $name = $ENV{incanickname} || $ENV{incareporter};
my $log =
    "Test '$ENV{incacomparison}' on output of series $name on " .
    "$ENV{incaresource} now produces '$ENV{incacomparisonResult}'";
my @timeBits = localtime;
my $timestamp =
    sprintf('%02d:%02d:%02d ', $timeBits[2], $timeBits[1], $timeBits[0]);

my $path = abs_path($ARGV[0]);
open(LOG, ">>$path") || die "Unable to open log $path";
print LOG "$timestamp" if defined($timestamp);
print LOG "$log\n";
close(LOG);
```

Report Filter Clarifies Reports

- Cooperates with Consumer for more accurate display
- Java class in Depot CLASSPATH extends `edu.sdsc.inca.depot.util.ReportFilter`
- `set/get` methods for `context`, `report`, `resource`, `stderr`, `sysusage`
- Invoked by Depot before incoming report is examined or stored
- Registered by setting `inca.depot.reportFilter` property in `$INCA_DIST/etc/common/inca.properties`

Depot Redundancy



Depot Redundancy

- Warm backup for the primary depot
- Synchronization with peers occurs at startup
- New configuration property

URIs of peer depots

```
inca.depot.peers=incas://some.other.peer:6324 incas://yet.another.peer:6324
```

Agenda -- Day 2

9:00 - 9:30	Inside the Inca Depot
9:30 - 10:10	Data display (data consumers)
10:20 - 11:00	Writing data consumers
11:00 - 12:00	Hands-on: Data display (data consumers)