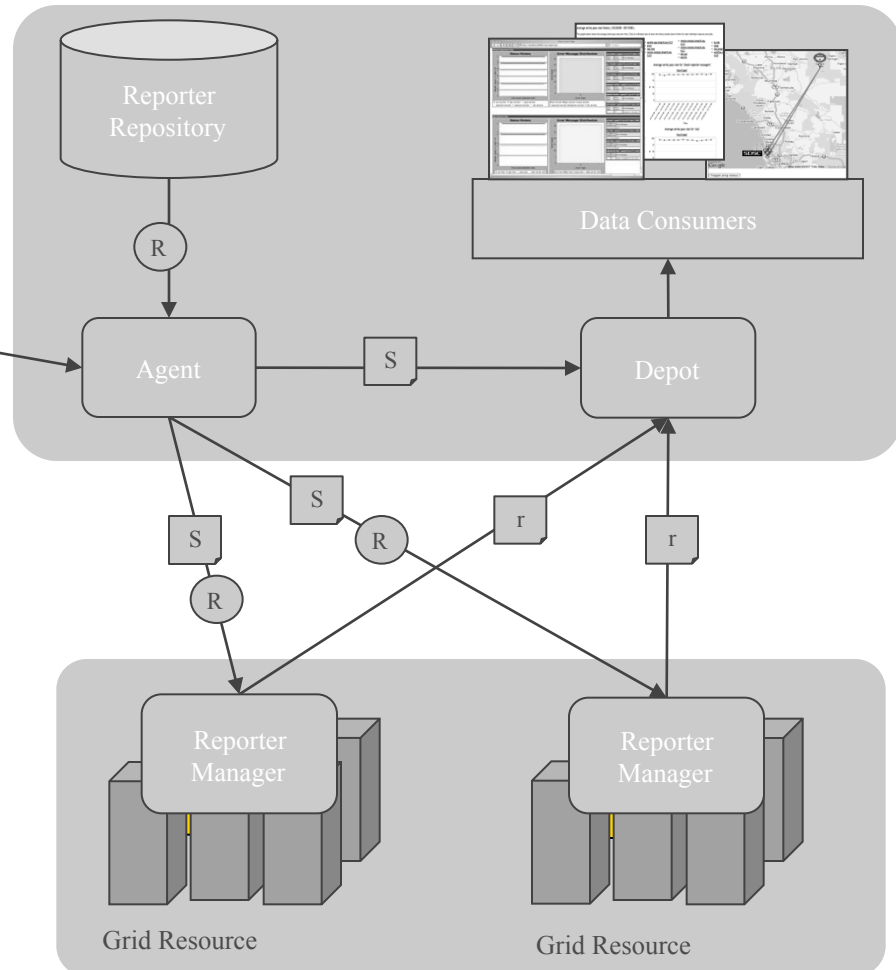
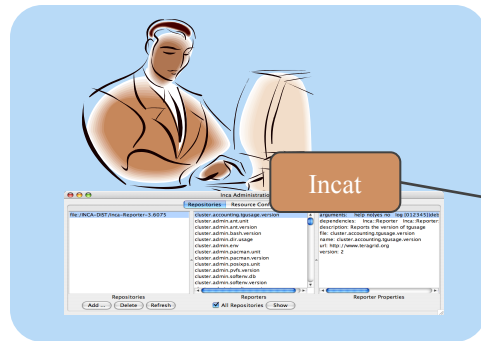

Administering Inca with incat

Shava Smallen
ssmallen@sdsc.edu

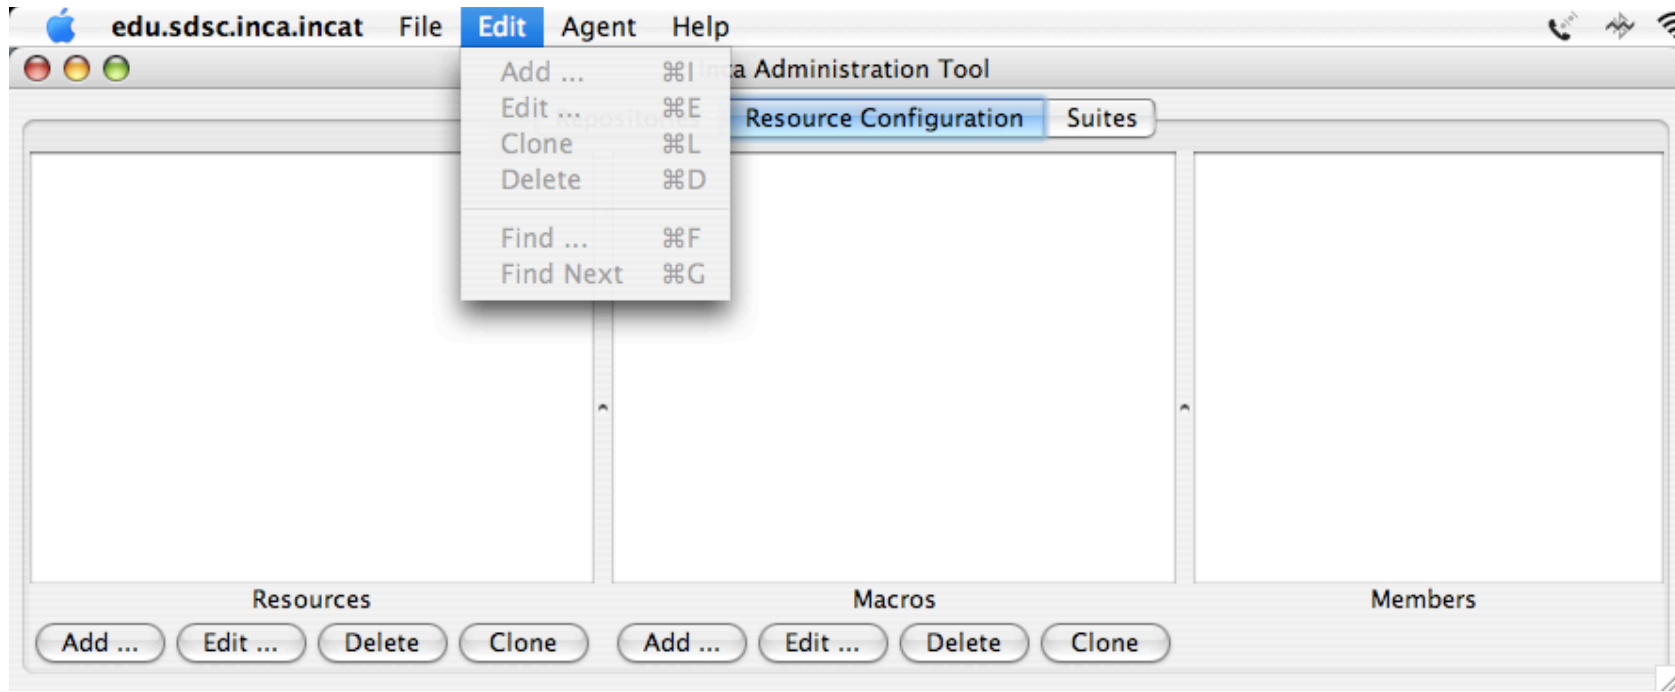
Inca Workshop
August 26, 2010

incat



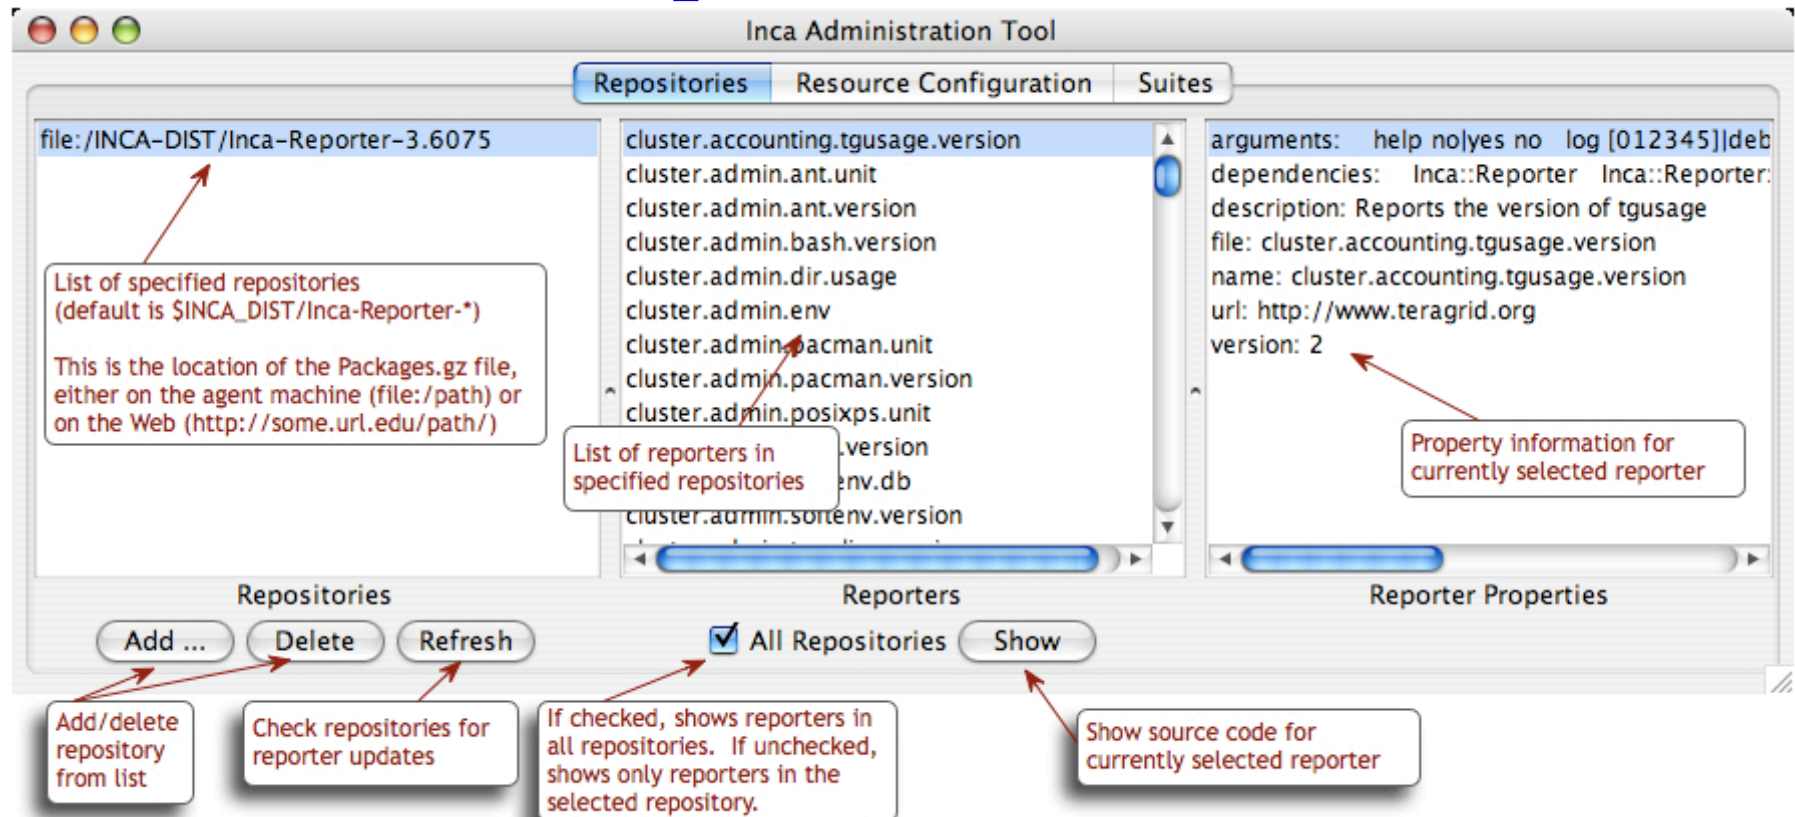
- A graphic editor for an Inca deployment configuration (XML describing repositories, resources, and series)
- Retrieves existing configuration from Agent and/or loads specified file

inca Display



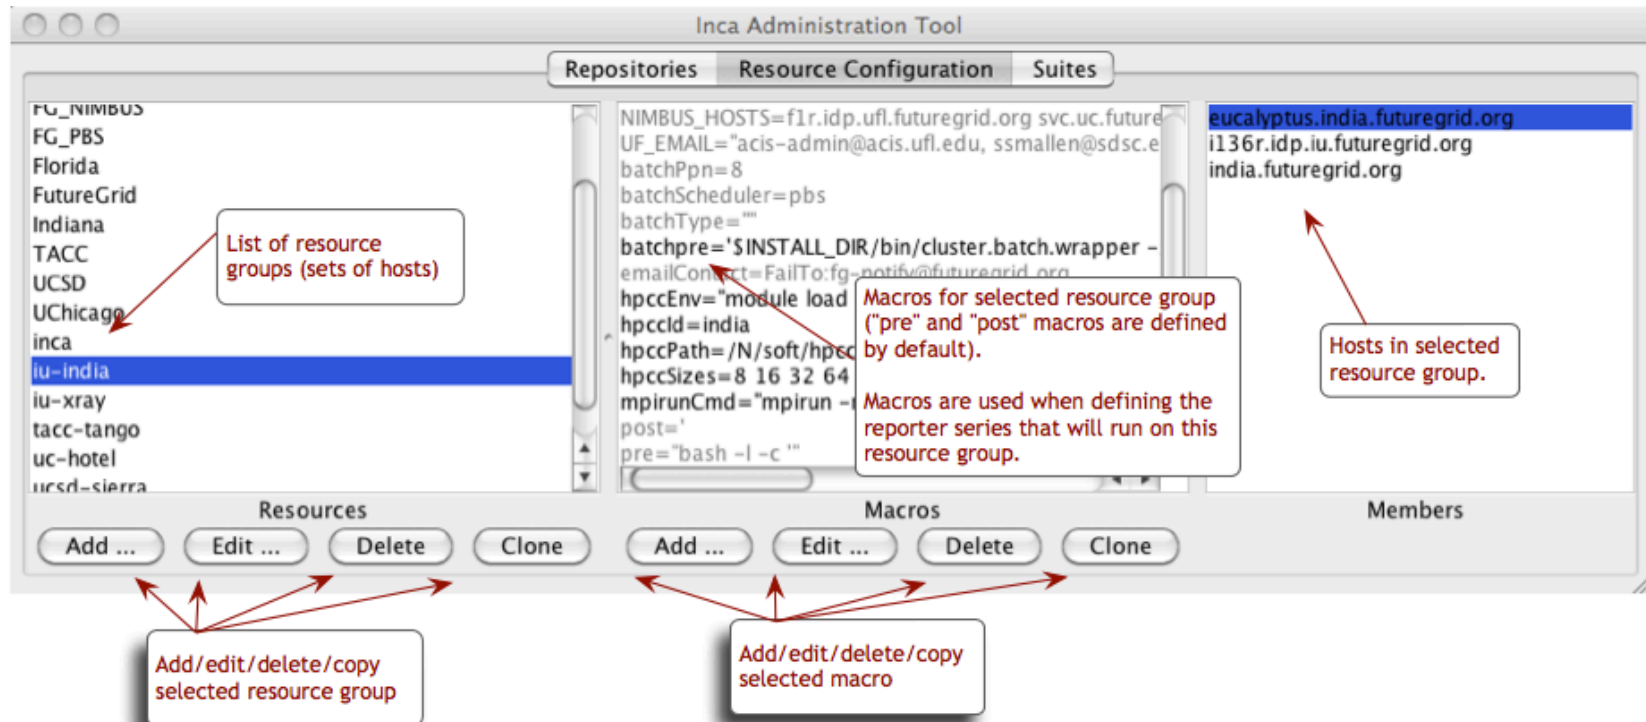
- Tabs for repositories, resources, and suites
- File Open/Save/Save As menu
- Edit menu supports modifying and searching selected panel
- Agent/Commit sends Inca deployment configuration to Agent
- Online user's guide via Help menu

Repositories Tab



- Repository list, reporter collection, catalog entry
- Reporter information retrieved from connected Agent

Resource Configuration Tab



- Resource name, associated macros, member hosts
- Inherited macros shown in gray
- Hosts with no IP address shown in red
- Resource “Add ...”, “Edit ...” open resource edit dialog

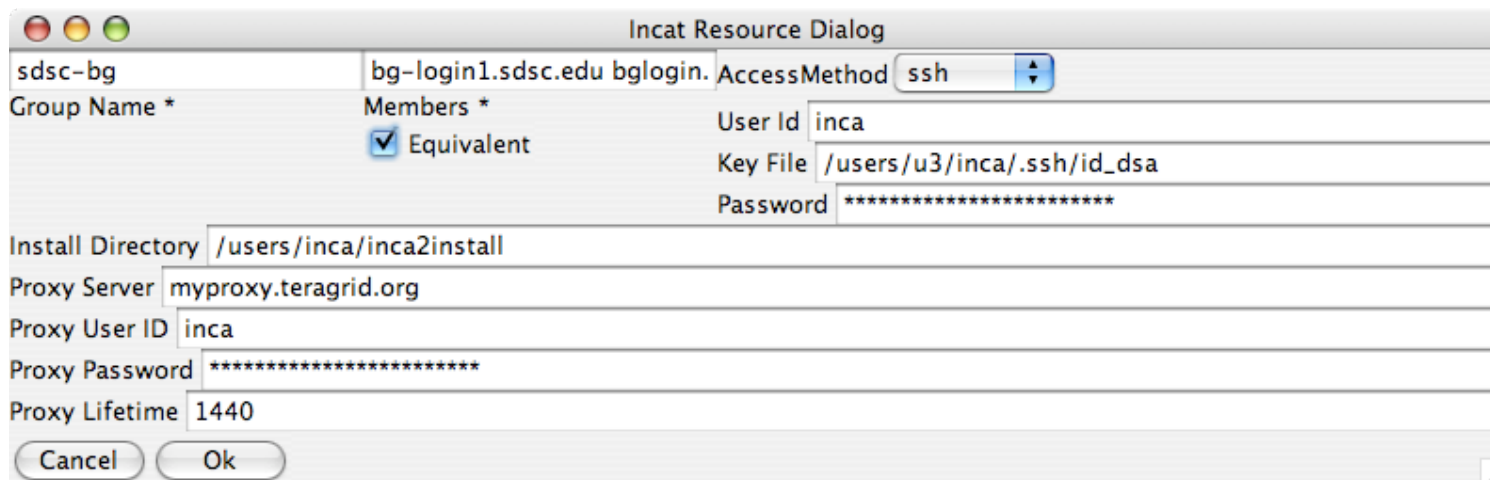
Resource Edit Dialog

The screenshot shows the 'Incat Resource Dialog' window with the following fields and callouts:

- localResource**: Arbitrary nickname for the resource group.
- client64-52.sdsc.edu**: Hostnames or other resource groups that are in this resource group.
- AccessMethod**: local (dropdown menu). Method the Agent uses to contact the Reporter Manager(s) in this group. Local, Globus2, Globus4, SSH, Manual.
- Group Name ***: localResource
- Members ***: client64-52.sdsc.edu
- Equivalent**: If checked, Inca will treat the hosts in this group as a single host, only running the reporters on one of them.
- Install Directory**: Directory where the Reporter Manager is installed. If blank, the Reporter Manager will be installed in `~/incaReporterManager`.
- Approval Email**: Resource administrator approves all changes.
- Proxy Server**: Proxy information for Agent access to Globus resources and for reporters that require a proxy.
- Proxy User ID**
- Proxy Password**
- Proxy Lifetime**

Buttons: Cancel, Ok

Specifying Ssh Resources



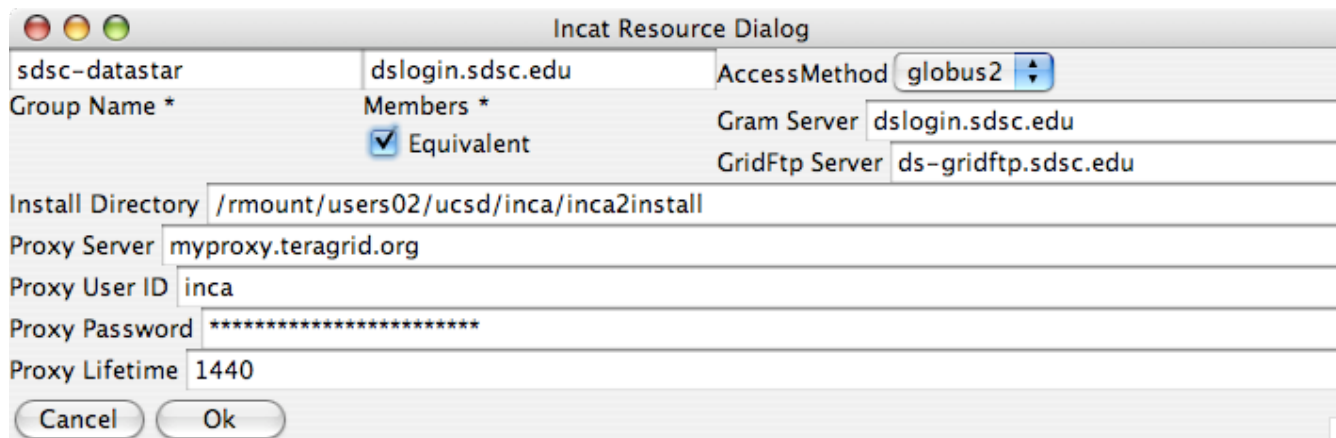
The screenshot shows the 'Incat Resource Dialog' window with the following fields and values:

sdsc-bg	bg-login1.sdsc.edu bglogin.	AccessMethod	ssh
Group Name *	Members *	User Id	inca
	<input checked="" type="checkbox"/> Equivalent	Key File	/users/u3/inca/.ssh/id_dsa
		Password	*****
Install Directory	/users/inca/inca2install		
Proxy Server	myproxy.teragrid.org		
Proxy User ID	inca		
Proxy Password	*****		
Proxy Lifetime	1440		

Buttons: Cancel, Ok

- Accessed by Agent using j2ssh classes
- Reporter Manager staging via sftp
- Passwords stored/transmitted encrypted

Specifying Globus Resources



The screenshot shows the 'Incat Resource Dialog' window with the following fields and values:

sdsc-datastar	dslogin.sdsc.edu	AccessMethod	globus2
Group Name *	Members *	Gram Server	dslogin.sdsc.edu
	<input checked="" type="checkbox"/> Equivalent	GridFtp Server	ds-gridftp.sdsc.edu
Install Directory	/rmount/users02/ucsd/inca/inca2install		
Proxy Server	myproxy.teragrid.org		
Proxy User ID	inca		
Proxy Password	*****		
Proxy Lifetime	1440		

Buttons: Cancel, Ok

- Accessed by Agent via Globus CoG Java classes
- Globus2 uses pre-WS APIs and Globus4 uses WS APIs
- Reporter Manager staged via GridFtp

Specifying Manual Resources

The screenshot shows a dialog box titled "Incat Resource Dialog". It has several input fields and a dropdown menu. The "AccessMethod" dropdown is set to "manual". The "Members" section has a checked checkbox for "Equivalent". The "Group Name" field is empty. The "Install Directory", "Approval Email", "Proxy Server", "Proxy User ID", "Proxy Password", and "Proxy Lifetime" fields are also empty. The "Cancel" and "Ok" buttons are at the bottom.

- Agent cannot use provided automated methods
- Resource administrator installs reporter manager manually and connects it to the agent

Grouping Resources

The screenshot shows a dialog box titled "Incat Resource Dialog". It has several input fields and a dropdown menu. The "Group Name *" field contains "compiler-gcc-fortran". The "Members *" field contains "jrdue-login PSC TACC NCAR". The "AccessMethod" dropdown menu is set to "globus2". There is an unchecked checkbox for "Equivalent". The "Gram Server" and "GridFtp Server" fields are empty. Below these are fields for "Install Directory", "Proxy Server", "Proxy User ID", "Proxy Password", and "Proxy Lifetime", all of which are empty. At the bottom of the dialog are "Cancel" and "Ok" buttons.

- Convenience for specifying series to run on multiple resources
- Directory, proxy info, etc. are like macros--most specific value is used for each host

Suites Tab

The screenshot shows the 'Inca Administration Tool' interface with the 'Suites' tab selected. The left sidebar contains 'Benchmarks', 'Cluster', 'Services', and 'Storage'. The main area is divided into 'Suites' and 'Series' sections. The 'Suites' section has buttons for 'Add ...', 'Edit ...', 'Delete', and 'Clone'. The 'Series' section has buttons for 'Add ...', 'Edit ...', 'Delete', 'Clone', 'Move To ...', and 'Run Now ...'. A list of series is displayed in the center, including 'aliasworks_to_@FG_ALIASES@', 'eucalyptus-webpage-loads_to_@EUCALYPT...', 'gcc_version', 'mpi_version', 'nimbus-external-telnet_to_@NIMBUS_HOSTS...', 'pbs_version', 'pgcc_version', 'ping_to_@FG_LOGINS@', 'ping_to_foxtrot.futuregrid.org', 'ssh_to_@FG_LOGINS@', 'ssh_to_foxtrot.futuregrid.org', and 'test_batch_job'. A 'Reporter Series' callout explains that a series is a set of reporters and reporter arguments, specifying the reporter, resource group, frequency, and arguments. Another callout explains that a suite is a set of reporter series. A third callout explains the 'Add/edit/delete/copy selected suite' button. A fourth callout explains the 'Add/edit/delete/copy selected series' button. A fifth callout explains the 'Move To ...' button. A sixth callout explains the 'Run Now ...' button.

Reporter Series
A series is a set of reporters and reporter arguments. Each series specifies a reporter to run, the resource group the reporter should run on, how often the reporter should run and what arguments the reporter should use when run.

Reporter Suites
A suite is a set of reporter series. Suites are a convenience for organizing related series. The Inca consumer uses suites for organizing the display of reporter results.

Add/edit/delete/copy selected suite

Add/edit/delete/copy selected series

Move selected series to a different suite

Execute selected series immediately on specified resource group and send results to depot

- Grouping of series into suites allows clearer Consumer display
- Report series named after reporter or given “nickname”

Series Edit Dialog

Incat Series Dialog

Reporter *

Resource *

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'

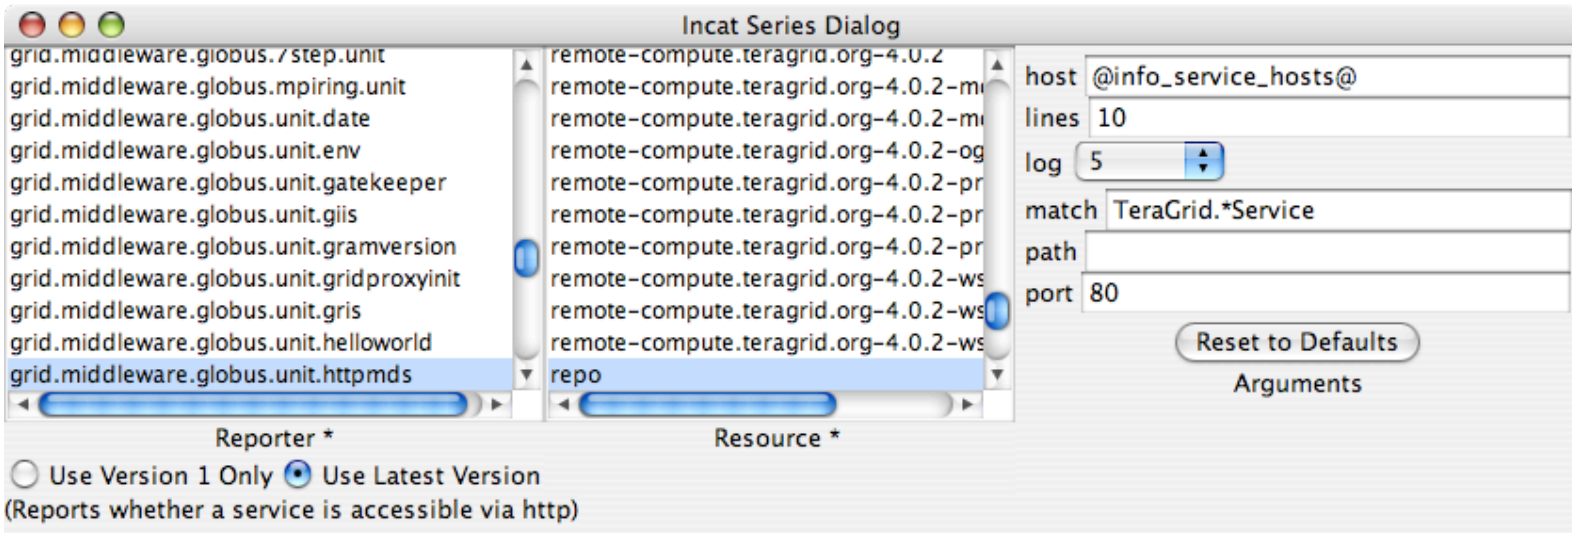
Notification Script EmailNotifier

Script Arguments inca@sdsc.edu

Cancel Ok

Reset to Defaults Arguments

Series Edit Dialog (Upper)



- Select reporter and resource; specify command-line arg values
- Argument list determined by reporter
- Use latest version automates Reporter updates
- Reporter description appears below reporter

Series Edit Dialog (Lower)

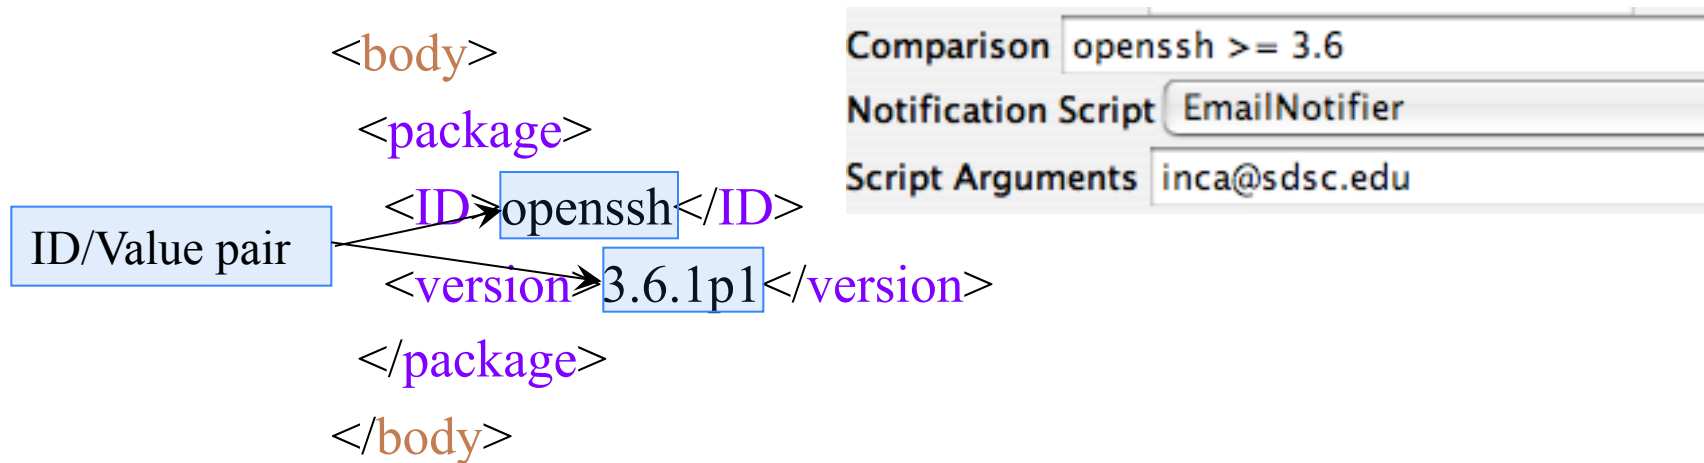
The screenshot shows a dialog box with the following fields and values:

- 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

Buttons: Cancel, Ok

- Run frequency uses extended cron syntax
- Nickname useful in later report display--should be unique
- Context allows specialized execution
- Reporter killed and error report stored if limits exceeded
- Comparison allows more precise definition of reporter success/failure
- Notification script run when state changes

Comparisons Refine Success and Failure



- Comparison can test body, errorMessage, ID values from report body
- Supports comparison operators and pattern matching
- EmailNotifier, LogNotifier provided; custom notifiers supported
- EmailNotifier reports success, failure, or both

Using Macros in Report Series

The screenshot displays the Inca Administration Tool interface. The 'Resource Configuration' tab is active, showing a list of resources on the left and configuration details for 'repo.teragrid.org' on the right. The configuration includes fields for 'cogClasspath', 'info_service_hosts', 'maxMan', 'maxSched', and 'maxWallTime'. The 'maxWallTime' field is highlighted with a red box. An 'Incat Series Dialog' window is overlaid on top, showing configuration for a series. The 'host' field is set to '@info_service_hosts@', and the 'Wall (Secs)' field is set to '@maxWallTime@'. Both macro references are highlighted with red boxes. Red lines connect these macro references in the dialog to their corresponding values in the resource configuration. The dialog also shows options for 'Reporter', 'Run Every', 'Nickname', 'Context', 'Limits', 'Comparison', 'Notification Script', and 'Script Arguments'.

- Separate series for each macro value
- Series with undefined macros references shown in red

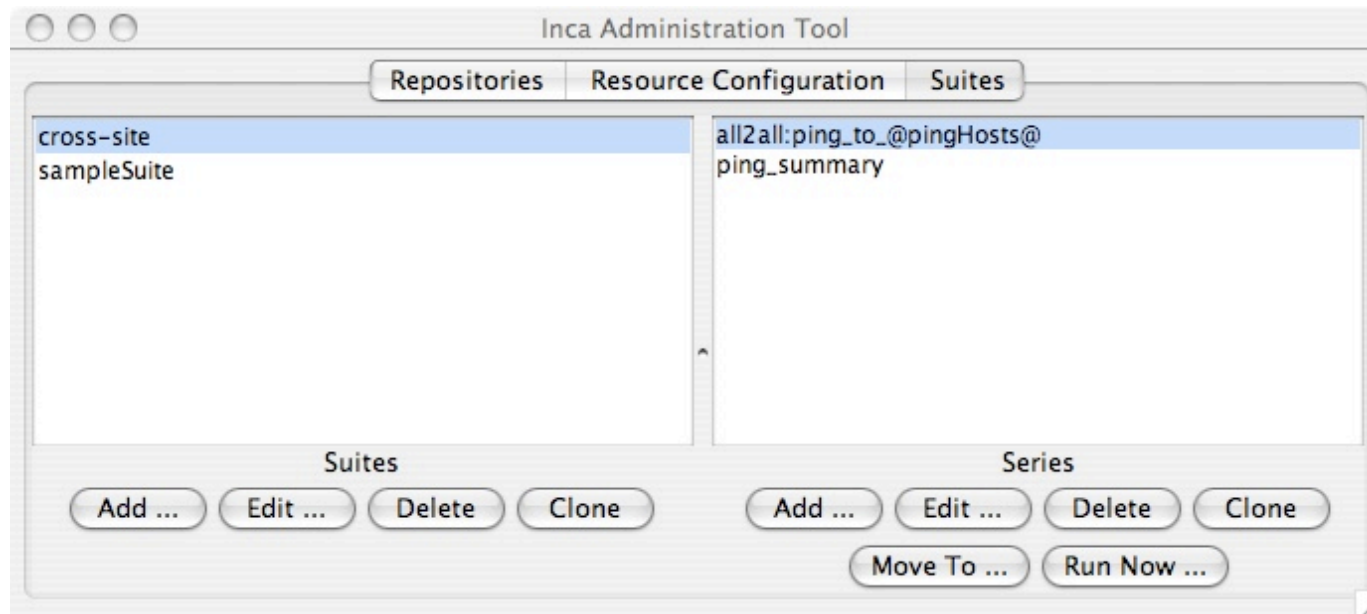
Executing series in batch mode

The screenshot shows the Inca Administration Tool interface. The 'Resource Configuration' tab is active, displaying a list of clusters on the left and their configurations in the center. The 'Incat Series Dialog' is open, showing configuration options for a series named 'hpcc@hpccSizes@'. Two callout boxes provide additional context:

- Define **batchPre** macro specifying batch wrapper, scheduler type, and special submit params**: This box points to the `batchpre` configuration line in the Resource Configuration window, which is `batchpre='$INSTALL_DIR/bin/cluster.batch.wrapper -scheduler=@batchScheduler@'`.
- Specify **batchPre** macro in context with extra batch wrapper arguments**: This box points to the `Context` field in the Incat Series Dialog, which contains `@batchpre@ -nodes=":@batchPpn@:@hpccSizes@" -walllimit=420 -exec='@hpccEnv@: performance.hpcc -args@post@'`.

The Incat Series Dialog also shows other configuration options such as `hpccPath`, `hpccSizes`, `mpirunCmd`, `post`, `Reporter`, `Run Every`, `Nickname`, `Limits CPU (Secs)`, `Wall (Secs)`, `Memory (MB)`, `Comparison`, `Notification Script`, and `Script Arguments`.

Executing all2all series



1. Specify an all2all series
2. Specify a summary reporter series using summary.successpct.performance
3. Add All2All filter in depot

Agenda -- Day 1

9:00 - 10:00	Inca overview
10:00 - 11:00	Working with Inca Reporters
11:15 - 12:00	Hands-on: Reporter API and Repository
1:00 - 2:00	Inca Control Infrastructure
2:00 - 3:00	Administering Inca with incat
3:15 - 4:00	Hands-on: Inca deployment