



LCG Security Group

Audit Requirements for LCG-1

<i>Date:</i>	19 June 2003
<i>Version:</i>	1.2
<i>Status:</i>	Draft
<i>Author:</i>	Ian Neilson

Document Log			
Issue	Date	Author	Comment
1.0	11 June 2003	Ian Neilson	Draft for Comment
1.1	18 June 2003	Ian Neilson	Some more details and add recommendations to GDB section
1.2	19 June 2003	Ian Neilson	Clarification of process accounting data volume

1 Introduction

This document deals with audit data retention at a Grid resource provider in order to map Grid Service Requests to the executables and identities that initiate them.

Unlike more traditional environments where the execution of a job on a UNIX system is most usually tied to an individual by user-id (UID) and group-id (GID) acquired at logon, the LCG Grid computing environment will derive the UID and GID local to the resource being used from a Distinguished Name (DN) contained in a digital certificate presented by the user on job submission. How the mapping from DN, which should uniquely identify an individual, to “resource-local” UID & GID and hence to executing process is accomplished may vary depending on site configuration.

This paper recommends the minimum set of data that should be retained by a grid resource provider in order to properly audit the execution of jobs at the resource. Resource providers should retain this information for a minimum of 90 days.

2 Audit Requirements

Job submission will normally progress from a User Interface (UI) machine, through a Resource Broker (RB) to a Computing Element (CE) and hence to the compute resource (usually a batch system). In some cases the RB is not used and the UI submits the job directly to the CE. Data access is through a Storage Element (SE) service.

2.1 User Interface

There are no requirements to retain data from UI machines.

2.2 Resource Broker

Since the RB is not necessarily local to the target resource it is not involved in the mapping of DN to UID. However it is only at this service that the origin (IP address) of the job submission is known and, as such, the RB service could provide valuable audit trail data following an incident.

The Security Group will make recommendations for audit data retention at the RB in the future.

2.3 Computing Element

The CE runs a process known as the *gatekeeper* which controls access to local resources through a mapping of certificate DN to local UID. Various mechanisms are in place or planned to allow or disallow jobs based on the accompanying certificate. Following a successful mapping to UID the *gatekeeper* creates a *jobmanager* process to manage the local submission and execution of the job (usually on the local batch farm). The *jobmanager* can be configured to query and output accounting data from the batch system on job completion.

The following paragraph applies to VDT release 1.1.8-7 and above.

The *gatekeeper* logfile (usually */var/log/globus-gatekeeper.log*) should be retained. By default, the output from the *jobmanager* appears in the *gatekeeper* logfile but the system can be configured for the *jobmanager* to log output to an alternative location. In this case both logs should be retained.

2.4 Storage Element

There is currently no SE deployment for LCG-1 but storage is accessed through *gridftp* services.

The *gridftp* service should be configured to log input and output transfers (*-i -o* options) with verbose logging (*-l -L* options) and the logfiles (usually */var/log/globus-xferlog*, */var/log/gsiwuftpd.log*) retained.

2.5 Batch System

Depending on local configuration and the batch system in use, the *jobmanager* may be configured to output relevant data to the *gatekeeper* logfile (see 2.3 above). This should include the following data for each job –

- the batch system job identifier
- the location (address/name) of the machine(s) used
- the time the job was started
- the time the job ended (or duration)
- the command executed.

In cases where the batch system data is not available to the *jobmanager*, the batch system logfiles should be retained separately and include the information listed.

Sufficient data should be retained from the batch farm nodes in order to trace the process activity resulting from a Grid Request. This data should include process accounting tables (Linux *pacct*) where this data is available. One estimate for compressed Linux process accounting data volume is 2.5Mb/day (225Mb total archive over 90 days) per node¹ on a reasonably busy machine. However, resource providers should clarify that this figure is sufficient for local resources and expected conditions.

3 Recommendations to GDB

The GDB is asked to approve that the following data be retained by **all LCG-1 sites** for a **minimum period of 90days**.

1. *Jobmanager* and/or *gatekeeper* logfiles (section 2.3 above).
2. Data transfer logs (section 2.4 above).
3. Batch system and process activity records (section 2.5 above).

¹ Estimated on a CERN batch cluster as the average of the 10 largest daily volumes seen in approx. 1 week from 354 nodes.