



LCG Grid Operations Service

LCG Service Level Agreement Guide



<i>Date:</i>	16 February 2004
<i>Version:</i>	0.3
<i>Status:</i>	Final
<i>Author:</i>	Trevor Daniels



Document Log			
Issue	Date	Author	Comment
0.1	28 Oct 2003	Trevor Daniels	First draft for discussion in GOC Steering Group and LCG Security Group
0.2	8 Dec 2003	Trevor Daniels	Includes comments from above.
0.3	16 Feb 2004	Trevor Daniels	General tidying up
0.3	3 June 2004	Ian Neilson	Moved to FINAL status following approval by GDB on 18 May 2004



1 Introduction

This document provides guidance to Resource and Service Administrators (hereafter called ‘Resource Administrators’) for meeting the requirements of the LCG Security and Availability Policy with respect to formulating and observing Service Level Agreements.

The adopted Security and Availability Policy states:

“The Administrator of each Service instance must maintain an assessment of the risks inherent in their particular Service design or resulting from local services or operational practice which might affect that Service’s Availability, Reliability or Performance, and publish the expected values of these service parameters in accordance with the GOC Procedures for Service Administrators. The publication of this information, together with other details described in the LCG Service Level Agreement Guide, constitutes the SLA with the user community for that Service.”

This document describes the nature of the Service Level Parameters for each LCG Service, the manner in which these parameters must be published, and the way in which the actual service level achieved is monitored by the Grid Operations Service¹. A companion document, the GOC Procedures for Resource Administrators, describes the way in which the target Service Level Parameters should be estimated by Resource Administrators.

2 Grid Services

The Grid Services to which Service Level Agreements apply are those which provide a production LCG service to a remote and general community of LCG users via the Internet. This definition would, for example, exclude a User Interface (UI) service restricted by a site firewall to a small local community, but would include a UI service intended for use by a remote user community via the wide-area Internet. Similar considerations of locality and access restrictions apply to all other Grid services.

Worker Nodes (WNs) are not individually subject to SLAs as there is no requirement for any individual WN to be available. The availability of the WNs as a service is included as one of the performance metrics of the Computing Element (CE) which feeds jobs to them.

The following Grid Services will be subject to SLA monitoring if they are published as providing a production quality service. This list may change over time.

CE (Computing Element) The Gatekeeper Service.

SI (Security Infrastructure) The several components of the security mechanism.

IS (Information Services) The GIIS interface publishing site information, the regional GIISes and the BDIIs associated with Resource Brokers.

LB (Logging and Bookkeeping) Accepts and logs events from the UI, RB and CE, and presents job status information to the user.

RB (Resource Broker) The Workload Management Service.

¹ In this document we use the term “Grid Operations Service” to mean the functions carried out by the LCG Grid Operations Centres acting in a coordinated fashion.



RC (Replica Catalog) Includes the Replica Manager, the Replication Metadata Catalog, the Replica Location Service and the Replica Optimisation Service.

SE (Storage Element) The interface to mass storage facilities.

UI (User Interface) Enables submission of a job to a RB. Requires an SLA only if its ssl port is externally visible.

3 Service Level Parameters

Each Grid Service is characterized by a number of Service Level Parameters. These are intentionally chosen to be conceptually simple and few in number, consistent with providing a useful indication to users of both the intended and actual quality of each service instance.

Three parameters are common to all services, but parameters describing Performance are specific to each Service.

3.1 Common Parameters

3.1.1 Scheduled Service Downtime

The Scheduled Service Downtime is a simple list of the dates and times when a Service is intended to be withdrawn from use and the expected duration of each withdrawal. Its purpose is to indicate when the Service will be unavailable to users for purposes of updating, maintenance or any other reason. Indicating the reason for the withdrawal is recommended but is at the discretion of the System Administrator. The information should be published as far in advance of the actual scheduled dates as possible.

The GOS will use this information in calculating the measured Availability and Reliability described below by discounting periods of Scheduled Service Downtime announced 24 hours or more in advance, but Service withdrawals, deliberate or otherwise, which occur with less than 24 hours' notice will be regarded as an unscheduled service failure.

3.1.2 Service Availability

Service Availability is defined to be the proportion of time that Service is (has been) effectively operational to the time it was scheduled to be operational expressed as a percentage, calculated over specified periods of operation. The time it was scheduled to be operational is wall-clock time minus Scheduled Downtime.

Three periods of operation are specified: Calendar Months, Quarters and Years. The target Service Availability will be assumed to be the same simple average for each of these (although they may be revised from time to time by the Service Administrator). The actual measured value will vary, and the historical records of measured values will be preserved by the GOS.

3.1.3 Service Reliability (MTTF)

Service Reliability is the reciprocal of the rate of Service Failures, where a Service Failure is defined to be one which results in an unscheduled break in that Service of sufficient magnitude and duration to be noticeable to users (defined for each Service), calculated over specified periods of operation. It is defined to be the Mean Time To Failure (MTTF) in units of days. The precise method of calculating Service Reliability is defined in the Appendix.

Three periods of operation are specified initially: Calendar Months, Quarters and Years. The target Service Reliability will be assumed to be the same for each of these periods (although they may be revised from time to time by the Service Administrator). The actual measured values will vary with time, and the historical records of measured values will be preserved by the GOC. It is expected that



measurements over Calendar Months will be discontinued when the Service Reliabilities routinely exceed 30 days.

3.2 Service-Specific Performance-Related Parameters

Service Level Parameters which indicate expected or actual performance are specific to each type of Service. It is anticipated that the particular performance metrics which are to be monitored will gradually evolve with time as operational experience with the Grid is gained.

3.2.1 CE

- a) Number of WNs available to this CE. Time average over calendar months.
- b) Turnaround time in each defined queue. Average over all jobs completed within each calendar month.

3.2.2 SI

- a) Time taken to obtain a certificate from a CA

3.2.3 IS

- a) Time taken to publish and retrieve specific information. Monthly average.

3.2.4 LB

- a) Time taken to retrieve logging information for a job

3.2.5 RB

- a) Time taken to complete job-list-match for specific test jobs

3.2.6 RC

- b) *{Performance metrics to be determined}*

3.2.7 SE

- a) Time taken for a monitor job running locally to the SE to retrieve a specific file.

3.2.8 UI

There are no performance metrics specified for a UI

3.3 Minimum Thresholds

Each of the Service Level Parameters has a defined threshold value which must be exceeded in quality for the Service to be acceptable as a production advertised LCG Service. These threshold values are determined each year by the GOC and the values for 2003 and 2004 are shown against each parameter in the Appendix. In later years, as the Services mature, these threshold values will be changed to reflect the higher service quality expected.

4 Target Service Level Parameters

The Service Administrator of an LCG Service is required to determine the target or design values of the Service Level Parameters of the Service (s)he is providing, and to publish them to the user community via the GOS website. The target value of a Service Level Parameter will be determined by



the design quality of the Service, for example the designed resilience of the essential components, the response of operational and systems personnel to problems, the availability of backup or hot standby components, etc. Performance will be determined by the scale of resources deployed for providing that Service and any other contention which there may be for those resources.

In summary, the Target value of a Service Level Parameter is the Service Administrator's estimate of the value (s)he expects the Service to achieve for that parameter.

Guidance on estimating the target Service Level Parameters from an assessment of the risks which might affect the Service and for estimating the performance of the Service is contained in the Resource Administrators' Guide.

5 Monitoring Service Level Parameters

The GOS will monitor the actual achieved value of all the defined Service Level Parameters for each Service instance by running well-publicised tests at agreed intervals. The precise tests and the intervals when they are executed are shown in the Appendix. *[This draft contains only a single example entry. The remainder are still under consideration.]*

The tests will be devised to minimize their dependency on extraneous variables as much as possible, and to be as straight-forward as possible.

Measuring actual Service Availability and Service Reliability will be carried out by testing for an appropriate response to some agreed Service-specific operation at an agreed interval; every 10 minutes is suggested. The test will result in either Success, Failure or be indeterminate (because, for example, the intervening network is down). The Availability is then estimated by the proportion of tests which result in Success over the sum of the Successes and Failures, excluding the periods when the Service is in Scheduled Downtime. Similarly, the MTTF is estimated from the number of transitions from Success to Failure over a certain period.

The Performance-related parameters are estimated in a similar way, by measuring the response to specified operations, although in this case the tests will take place less frequently, typically every hour.

Other parameters, such as the number of WNs available, may be determined by a simple extraction of information from the appropriate information server.

6 Publishing the Service Level Agreement

Each site which offers production LCG Services is required to publish a Service Level Agreement (SLA) in a prescribed format on the GOC website. The SLA is comprised of four parts: the Statement of Compliance, the Schedule of Downtime, the Target and Measured Service Level Parameters, and the Exceptions to the Security and Availability Policy.

6.1 Statement of Compliance

The LCG Security and Availability Policy requires each site offering LCG Services to conduct a self-audit of compliance with the provisions of that Policy at least every two years. This section of the SLA includes a summary of the findings of the most recent audit in a prescribed format together with a statement that

- a) the Service will continue to be operated in compliance with the LCG Security and Availability Policy and its associated documents (see <http://cern.ch/proj-lcg-security/documents.html>) ,
- b) LCG updates will be applied to the Service promptly, especially those which include security-related patches, and



- c) the site will make provision to respond to security incidents as required by the Agreement on Incident Response.

6.2 Schedule of Downtime

The Schedule of Downtime is posted by the Resource Administrator, and consists of a list showing the times when each of the Services will be withdrawn and the times they will be returned to service, together with the reason for the Service withdrawal, with later times appearing first. All times will be shown in both local and UTC (GMT) bases. The time of posting the information will be automatically added. An existing posting can be revised by posting an updated version of the entry, when both versions will remain visible with the superseded one so marked.

6.3 Service Level Parameters

This section includes Threshold, Target and Measured values of all Service Level Parameters appropriate to each of the Services being offered.

Threshold values are posted by the GOS, with the current values being immediately visible and previous values available for viewing in a history file.

Target values are posted by the Resource Administrator, with the current values being immediately visible, together with the date of posting, and with earlier postings available for viewing in a history file.

Measured values are posted by the GOS as they become available, and are presented in a way that enables direct comparison with the Target values.

Graphical presentations of all these data will be available.


6.4 Exceptions to the Security and Availability Policy

The Security and Availability Policy makes provision for country-specific exceptions or extensions to be made to the Policy if they are required in order to comply with that country's legislation. If a production Service is affected in any way by the exercising of this provision it is required that the exception be described in this part of the SLA for that Service, with the reason for the exception or extension clearly stated. These exceptions are posted by the appropriate Service Administrator.

Appendix

Calculating Service Reliability (MTTF)

Define the period over which the MTTF is to be calculated as starting at time t_2 and ending at time t_3 . Locate the latest failure earlier than t_2 and find the time t_1 at which it occurred, and locate the earliest failure later than t_3 and find the time t_4 at which it occurred, ie the two failures nearest to but outside the period in question. Count the number of failures, including these two and all failures in between them. The MTTF for the period t_2 to t_3 is then given by $(t_4 - t_1) / (n - 1)$.



Note that this definition works even when there are no failures within the defined period, and gives a suitably indicative value (see the second example) for any period.

Examples. Failures occurred on 2 Mar, 5 Apr, 10 Apr, 28 Apr, 4 Jun, 28 Jul.

The MTTF for Apr is given by:

first failure: $t_1 = 2$ Mar, last failure $t_4 = 4$ Jun, so $t_4 - t_1 = 63$ days.

$n = 5$, so MTTF for Apr is $63 / 4 = 15.75$ days.

The MTTF for May is given by:

first failure: $t_1 = 28$ Apr, last failure, $t_4 = 4$ Jun, so $t_4 - t_1 = 37$ days

$n = 2$, so MTTF for May is $37 / 1 = 37$ days

The MTTF for Jun is given by:

first failure: $t_1 = 28$ Apr, last failure, $t_4 = 28$ Jul, so $t_4 - t_1 = 61$ days

$n = 3$, so MTTF for Jun is $61 / 2 = 30.5$ days

The MTTF for Q2 is given by:

first failure, $t_1 = 2$ Mar, last failure, $t_4 = 28$ Jul, so $t_4 - t_1 = 148$ days

$n = 6$, so MTTF for Q2 is $148 / 5 = 29.6$ days

Note that a lower estimate can be made of the MTTF of any period as soon as that period has ended, since n is then known, and t_4 must be greater than the end of the period. The MTTF for that period to date t up until the next failure can then be expressed as " $> (t - t_1) / (n - 1)$ days".



Service-Specific Details

{Only the first row of this table is filled in and the values shown are suggestions for discussion before attempting to complete it}

Service	Parameter	Threshold	Test Frequency	Test
CE	Scheduled Downtime	2003: 10 hrs/mth 2004: 20 hrs/yr	n/a	As published
	Availability (quarterly)	2003: 90% 2004: 95%	10 mins	ce-authentication
	Reliability (quarterly)	2003: 15 days 2004: 40 days	10 mins	ce-authentication
	# WNs	2	1 day	As published
	Turnaround (Short)	30 mins	1 hr	Job submit
	Turnaround (Long)	24 hours	n/a	Stats
	Turnaround (Infinite)	7 days	n/a	Stats
SI				
IS				
LB				
RB				
RC				
SE				
UI				