

# Logging Proposal for JRA1 Code

*JRA1 All-Hands, Helsinki, 18-20 Jun, 2007*

*Steve Fisher/RAL and Ales Krenek/Cesnet*

[www.eu-egee.org](http://www.eu-egee.org)  
[www.glite.org](http://www.glite.org)

- **Why the logging discussion – again?**
- **Aims**
- **Specific recommendations**
- **Next steps**

- **Have discussed on several occasions:**
  - Each time the answer was to adopt some kind of logging framework
- **Re-triggered by David Groep's document: [Middleware Security Audit Logging Guidelines \(version 0.8\)](#)**
  - This document says what should be logged (anything security related) and how it should be logged (via syslog)
  - It is also the basis of some ad-hoc OGF agreement with OSG
- **We were asked to produce a recommendation**
  - This was sent last week to all JRA1 members - so far it has been remarkably quiet.
  - Aim to produce final version next week

- **David's document matches its title.**
  - He is only concerned about security audits.
- **Logging serves multiple purposes:**
  - developers to check program logic
  - developers to understand problems on a running system
  - sys-admins to check for correct functioning
  - sys-admins to investigate possible security incidents

- **Unify configuration of logging**
  - This is the main requirement, it is felt to be unavoidable for smooth deployment of the services
  - Other requirements follow from it
- **Follow the spirit of the security audit logging guidelines document**
  - This is a good document with valid concerns
- **Uniformity of what is logged**
  - makes life easier for site administrators
  - improves communication between developers when interaction between middleware components is analysed.
- **Minimise number of solutions**
  - With diverse developer communities we are unlikely to find a solution which suits everybody

- **Sys-admins determine what is logged at their site**
- **If sys-admin suspects a problem he may want to increase the logging**
- **Should be easy to do this without intimate knowledge of the middleware**

**All JRA1 components MUST use a common configuration file format modeled on the log4j one ... (The configuration file should be in XML as it is more powerful than a properties file).**

- **Java**
  - either log4j or Jakarta commons
- **C++**
  - Log4cxx
  - Should try HEAD version and perhaps contact developers
- **C**
  - log4c
- **Python**
  - log4py

**A dedicated *security* logger MUST be used for security related messages**

**A dedicated *access* logger MUST be used to log initial access to a service**

**A dedicated *control* logger MUST be used for service startup, configuration and shutdown**

**Other loggers SHOULD be named in a shallow hierarchy - not simply one logger per class.**

- **Finer severity levels (e.g. as defined by syslog) may be used when appropriate.**
- **If you use tomcat this handles the access log**



- **Brief but include as much relevant information as practical.**
- **Only log at most 768 bytes**
  - Some syslog implementations cannot deal with messages of more than 1024 bytes,
  - Restriction only applies to the loggers control, security and access to ensure that they can be sent to syslog.
- **Any key-value pairs should be separated by a comma (which is preferred) or a space; and the key and value should be separated by an "="**
- **White space should be used as necessary for readability**
- **Only characters from the seven-bit ASCII set with decimal values in the range 32-126 should be used in the message.**
  - It can be assumed that the syslog appender can translate or escape anything nasty.
- **Messages must be self-contained.**
  - If your message is too long devise a way to split it into multiple stand-alone messages

## **Service start up, configuration and termination**

To control logger

## **Authorization logging**

To security logger

## **Initial service access, session establishment and termination**

To access logger

- **Other Logging**
  - as the developer sees fit

- **If the request processing etc. is related to an externally identifiable entity (jobid, file guid, etc.) this identifier should be logged.**
- **Whenever there is a logical notion of session, request, etc., and multiple messages related to one such entity instance can be logged, it is recommended that a unique id is assigned to such an entity and it is logged in all the messages consistently.**

- **We hope we can agree to go forward with this proposal today**
  - Perhaps with minor changes
- **Then need a timescale for implementation**
  - Otherwise, once more, nothing will actually happen

# Comments ???