Document: GWD-Perf-14-1
Category: Informational
Obsoletes:
Expires:January, 2002

Dan Gunter
Lawrence Berkeley National Lab
Brian Tierney
Lawrence Berkeley National Lab
Ruth Aydt
University of Illinois, Urbana-Champaign
July 2001

Timestamp Model for Grid Computing

1.0 Motivation

There are many occasions for producing and consuming timestamps in computing, ranging from performance analysis to security protocols. Grid applications and protocols which produce and consume timestamps need to interoperate, either in real-time or by storing and using information in directory services or archives. A standard model for timestamps will allow standard formats to be constructed, and provide a basis for translating between formats.

2.0 Scope

This document proposes a *model* for timestamps. It does not make recommendations on how to gather and verify the accuracy of underlying timestamp values. At the time of this writing, a *best practices* document for creating timestamps does not exist, although this would be a valuable addition to the Grid Forum deliverables.

This document also does not attempt to model any other metric of time aside from a single moment. More complex constructions involving time, such as time intervals or time series, should be discussed in a separate document.

Finally, the timestamps presented here are primarily intended to represent real measured times, such as those returned from a host's clock, or a time logged into a system log or web log.

3.0 Model

A timestamp is an estimate of a single moment in time. In most applications, the timestamp describes when a real-world event occurred, and thus the timestamp is only an estimate of the "real" time of the event. In order for this estimate to be used in analyses, it should carry with it an estimate of its own *resolution* and *accuracy.* Thus, the timestamp has three parts: time value, resolution, and accuracy. All three parts of a timestamp are required, although timestamp formats may use special "null" values to indicate that the information is not available.

3.1 Time Zones

The time of day which most people are familiar with is really composed of two things: a time, and a time zone. The time zone is the offset of the time from Universal Coordinated Time (UTC), which is the more precise standard replacing the perhaps more familiar Greenwich Mean Time (GMT).

In order to avoid the overhead of adding and subtracting time zones from timestamps in the intermediate processing stages, timestamps for the Grid should always be in UTC time.

3.2 Timestamp Value

The timestamp value is measured in seconds and fractions of a second UTC. Calendar dates up to

the limit of 4 digits (year 9999) at picosecond resolution should be possible.

3.3 Timestamp resolution

The definition of *resolution* used here is identical to that given in [1]:

> A clock's "resolution" is the smallest unit by which the clock's time is updated. It gives a lower bound on the clock's uncertainty. (Note that clocks can have very fine resolutions and yet be wildly inaccurate.) Resolution is defined in terms of seconds. However, resolution is relative to the clock's reported time and not to true time, so for example a resolution of 10 ms only means that the clock updates its notion of time in 0.01 second increments, not that this is the true amount of time between updates.

The representable range must be at least 10^{-10} seconds (picoseconds) to 10^{3} seconds (~15 minutes). The upper end of this range is somewhat arbitrary, and it is suspected that almost all resolutions will be sub-second. The resolution is often determined by the underlying operating system and the programming language being used. For example, the standard timing call in C returns a resolution of microseconds, where Java, Perl, and Python all return a value with a resolution of milliseconds. Some operating systems also provide access to a clock with nanosecond resolution. (e.g., the Solaris and RTLinux gethrtime() call).

3.4 Timestamp Accuracy

The *accuracy* of a timestamp summarizes the how close the reported value is to the “true” value, or the “margin of error” for the time value. As noted above, this value varies independently from resolution. The accuracy of a timestamp is represented in seconds, with a range of at least picoseconds (as for the timestamp value) up to a day (24 hours).

4.0 Summary

This document has presented a simple model for representing timestamps. Hopefully, the existence of a separate model will make generation of various timestamp formats quicker and less error-prone.

5.0 References

[1] V. Paxson, G. Almes, J. Mahdavi, M. Mathis, “Framework for IP Performance Metrics”, RFC 2230, May 1998