Uses of network monitoring information within the EDG

P. D. Mealor

March 25, 2003

1 Version history

24 March 2003 Initial version

2 Introduction

This document describes the uses of network monitoring information within the EDG. Monitoring information is published using the R-GMA information system.

3 Network cost suite

The main use of monitoring information published via R-GMA is to calculate what is known as the network cost function. The network cost function is a calculation of the "cost" of a file transfer for a file of a particular size between pairs of Storage Elements. This cost is effectively the time it will take to perform that transfer, and has an associated uncertainty. The network cost function is calculated by the Network Cost Suite[3].

The cost is used to optimise data replication within the datagrid. ¡¡more¿¿

The network cost function in its current form uses, simply, TCP achieved throughput information. For EDG, this information is measured by the Iperf tool and the IperfER infrastructure for pairs of network monitor hosts.

The information is matched to pairs of Storage Elements by using the linking tables described in section §A.2. The network cost can then be calculated by simply dividing the file size by the achieved TCP throughput.

The network cost function has two methods of operation. In the first, a single pair of Storage Elements is specified, and the network cost function is calculated for that pair. In the second, a list of sources and a list of destinations are specified. The network cost function is then calculated for each source and destination pair and a matrix of costs is returned.

3.1 Future work

Currently foreseen extensions to the Network Cost Suite are:

- A version of the suite which can extract information from a Web Service.

- Implementation of different cost models, so as to be able to calculate the network cost for interactive or other non-bulk data transfers. This calls for the use of characteristics other than simple achieved throughput.

- Implementation of an OGSA[2, 1] compliant Grid service.

A Information model

Note that a more complete explanation of the schema and information system is given in [4].

A.1 Network monitoring information

Network monitoring information is split into tables according to the metric being measured. Every table has the following columns: source, destination and timestamp. The source and destination are the hostnames of the source and destination of a measurement, and the tool used to make the measurement. Most tables also contain either a value column or minimum, maximum and average columns, depending on the measurement: a throughput just has a value, while RTT has a min, max and average.

Each table also has columns containing the values of parameters of each measurement. Parameters include such things as the packet sizes used, buffer sizes and so on.

A.2 Linking information

There are two additional tables used to link Storage Elements and Computing Elements to the network monitoring information. These tables just contain a SE or CE identifier and the name of a network monitoring host nearby. The network monitoring host(s) is chosen in the expectation that measurements made to or from it will closely match a similar measurement made to the SE or CE.

References

[1] *Open Grid Services Infrastructure (OGSI).* http://www.gridforum.org. Editors: S. Tuecke and K. Czajkowski and I. Foster and J. Frey and S. Graham and C. Kesselman and D. Snelling and P. Vanderbilt.

[2] Ian Foster, Carl Kesselman, Jeffrey M. Nick, and Steven Tuecke. *The Physiology of the Grid.* http://www.globus.org/research/papers/ogsa.pdf.

[3] Robert Harakaly and Franck Bonnassieux. *Network Cost Suite programmer's manual.*

[4] P. D. Mealor. *Description of the EDG network monitoring schema for R-GMA.*