Enabling Network Measurement Portability Through a Hierarchy of Characteristics

Bruce B. Lowekamp
William and Mary
lowekamp@cs.wm.edu

Richard Hughes-Jones
University of Manchester
R.Hughes-Jones@man.ac.uk

Brian Tierney
LBNL
bltierney@lbl.gov

Thilo Kielmann
Vrije Universiteit
kielmann@cs.vu.nl

Les Cottrell
SLAC
cottrell@slac.stanford.edu

Martin Swany
University of Delaware
swany@cis.udel.edu

Abstract

Measurement and prediction of network resources are crucial so that adaptive applications can make use of Grid environments. Although a large number of systems and tools have been developed to provide such measurement services, the diversity of Grid resources and lack of central control prevent the development of a single monitoring system that can be deployed to answer every application's resource queries for connections between any pair of machines it can use. We propose a standard for representing network entities and measurements of their properties. Our standard enables the exchange of measurements and will allow applications to function even in environments without the particular measurement system for which they were developed. We present an overview of our measurement representation and evaluate its usefulness. We have used the characteristics hierarchy to store and exchange measurement data between several systems, and we discuss its usefulness in comparing the output of several measurement tools.

1. Introduction

One of the defining characteristics of Grid environments is the constant change in communication and computational resources. To deal with the dynamic nature of these resources, measurements of current state can be used to drive resource allocation and scheduling decisions. For example, network monitoring data can be used by Grid data management middleware when selecting the best source from which to copy replicated data, or by resource brokers when assigning resources to a particular application. Providing these services requires historical data and, where possible, predictions of the future behavior of the available network resources. Network monitoring systems that record network conditions before and during applications are an important component of the Grid tool chain, and many have been developed to serve a variety of purposes [4, 7, 13, 16, 25, 26]. Between all of the available monitoring systems, few users will have questions about network capabilities that cannot be answered.

As an example, a Grid Information Service (GIS) may publish achievable throughput information from an iperf probe. Using this achievable throughput information, higher-level Grid services can be built to answer queries such as: "How long will it take to transfer 5 GBytes from Host A to Host B using GridFTP with 4 parallel streams?" The answer to this question might be found by applying a prediction algorithm (e.g. [22]) to results from previous transfers that happened at roughly the same time of day, on the same day of the week, etc. Grid Scheduling Services and Grid Replica Management Services could use such a service to collect data to guide scheduling and data placement decisions.

While some Grid applications have very specific questions to ask a GIS, general adaptive Grid applications have two basic requirements:

- They must have a measurement between each pair of machines the application could run on.

- The measurements between different pairs must be comparable.

A single monitoring system deployed across the entire Grid could provide such information. However, the reality of grid environments is somewhat different:

- Each system monitors only the parts of the network where it has been deployed, configured, and/or authorized to run.

- Different systems measure different properties, utilize different probes for the same property, or use the same probes with different parameters, according to the interests of their managers.

The first point is easily recognized. The second is more subtle—after all, what most applications want to know is the bandwidth between two machines. But the subtleties of *what* is being measured makes even such a simple question difficult. Users and applications generally want to know how much bandwidth is *available* on the path. Network engineers and protocol tuners generally want to know the bottleneck *capacity* along a path. Both groups refer to what they measure as *bandwidth*.

Agreeing on a single, universal Grid measurement system and a small set of universally-used metrics would be one approach to these problems. However, the natural strides of research will bring the development of multiple network measurement systems and newer measurement tools, and groups will form virtual organizations that select their own measurement system. Even when the same system is selected, it may be deployed using probes with different parameters.

Although, or perhaps because, the authors have contributed their share of incompatible network monitoring systems, our goal is to develop techniques to promote compatibility between monitoring systems, so that an application can make use of data gathered by different systems using different probes. We face two primary challenges:

- A uniform way of representing network topology and structure, so that measurements can be described by more than just their endpoints

- A way of recognizing and comparing measurements of the same network property

We address both issues in this paper. Resolving the first is a matter of developing a portable way to express network topology that encompasses the representations used by existing monitoring systems. The second problem is more challenging. In this paper, we focus on naming the network properties being measured as the first step toward achieving the full goal of comparing disparate measurements. The application of this naming system will facilitate the creation of common schemas for describing network monitoring data in Grid Monitoring and Discovery Services, and thus help address portability issues between the wide variety of network measurements used between sites of a Grid.

Many participants in the Network Measurements Working Group (NMWG) of the Global Grid Forum have contributed to the characteristics hierarchy presented in this paper [15]. The NMWG is engaged in the ongoing task of developing standard schemas using this characteristics hierarchy to achieve the goal of interoperability between different
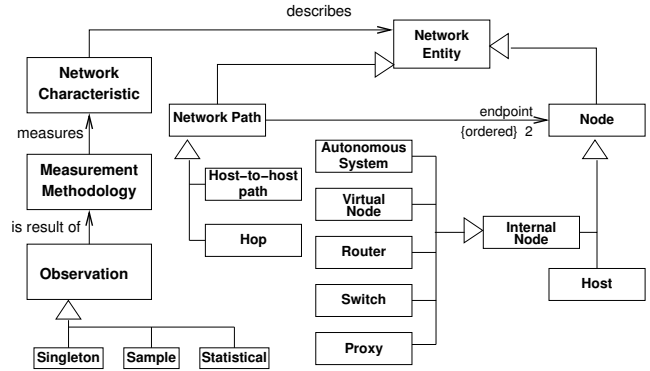

Figure 1: The relationship between observations, measurement methodologies, characteristics, and the various types of network entities.

monitoring systems. In this paper, we present a summary of the NMWG's work on measurement characteristics and an evaluation of how the characteristics ensure measurement portability. Our goal in this paper is to demonstrate that the characteristics we describe are sufficient, in terms of breadth and depth, to enable the exchange of network measurements currently taken by Grid monitoring systems.

Section 2 discusses the terminology used to represent network measurements. Section 3 describes how network topology is represented. Section 4 summarizes the network characteristics that are commonly measured. Section 5 evaluates how well the proposed characteristics hierarchy can be used to enable measurement portability.

2. Representing Measurements

There are two principle pieces of information in a network measurement:

- the *characteristic*, or property, being measured

- the *network entity*—path, link, switch, etc.—whose characteristic is being measured

Beyond those two, essential information includes the *methodology* used to measure the characteristic and the conditions under which each measurement (each *observation*) was taken. These factors are represented in the UML diagram shown in Figure 1. Boxed arrows (→) indicate an inheritance, or "is-a" relationship, and open arrows (→) indicate attributes. This section discusses how these factors come together to represent a network measurement.

2.1. Characteristics and Methodology

A characteristic is an intrinsic property that is related to the performance and reliability of a network entity. More

specifically, a characteristic is a primary property of the network or of the traffic on it. A characteristic is the property itself, not an observation of that property. Example characteristics include capacity and one-way delay.

There may be multiple ways to measure a characteristic. A measurement methodology is a particular technique for recording or estimating a characteristic. Different tools may implement the same methodology, or a single tool may offer a choice of methodologies.

2.2. Network Entities

Network observations can be reported for a variety of different types of network entities:

End-to-End Paths are the common case of host-to-host measurements. In Figure 1 we refer to them as *host-to-host paths*.

Links between routers or switches are measured for capacity, availability, and latency. In Figure 1 we refer to them as *hops*.

Nodes may report useful information such as router queue discipline or host interface speed. In Figure 1 we distinguish between *hosts* as our measurement endpoints and *internal nodes*, found along a path in between two *host* nodes.

Furthermore, to systems that report network topology, the *existence* of the network entities themselves—hosts, routers, switches, and links—form the actual observations being reported.

Keeping in mind our goal of developing a representation flexible to be used by all monitoring systems, we adopt an entity representation consisting only of *Network Paths* and *Nodes*, as shown in Figure 1. While some representations focus more on distinguishing between paths, links, and hops, we avoid making that distinction at the first level of our representation to escape imposing structure that would be incompatible with some systems. All kinds of paths and links are of type *path*, and the specific type—end-to-end, layer 3, and layer 2—can be distinguished by the endpoints. A measurement system that distinguishes between paths, hops, and links can select its most appropriate internal representation with this information. The topology representation is discussed more thoroughly in Section 3.

2.3. Attributes

Networks can differentiate their behavior based on a wide range of traffic characteristics. For example:

- Choice of protocol (e.g. TCP vs. UDP) can influence a network's behavior.

- Different Quality of Service (QoS) levels affect all aspects of network behavior; in fact, some QoS policies may specify using completely different routes between the end hosts.

- Route flaps or other instabilities mean that the same end-to-end traffic may experience a completely different environment from moment to moment.

- In high-bandwidth environments, the hosts performing the measurement may be the bottleneck, rather than the network path.

To represent these qualities, nodes and paths are annotated with attributes such as the protocol stack (e.g., TCP over IPv4) and QoS level. These attributes describe the conditions under which a particular observation was made, and a carefully designed monitoring system can make use of this feature to report the behavior (including routing) of a network for a variety of traffic types. The attributes annotating a node or path are not merely descriptive, but form a tuple that defines a unique node or path. In this way, the model allows superimposition of multiple topologies over the actual physical nodes and paths, allowing the behavior of the network under different types of traffic to be characterized separately.

Very few attributes are truly mandatory; for example, a system collecting regular iperf measurements frequently will not take corresponding traceroute observations, although the results of iperf would be significantly affected by route changes. However, the goal is to support the representation of more precise measurements when such information is available. In each specific case, it is the task of the monitoring system to provide sufficiently precise attributes.

3. Topology Representation

Beyond the representation of nodes and edges in topology graphs, a topology system for measurement interchange must address the challenge that there are different approaches to representing topology, ranging from the extremely physically accurate approach, via representing only approximate grouping or closeness of sites, to only storing the host pairs used in end-to-end measurements. Here we describe how our proposed system addresses these issues. Topology is represented using both the entity types shown in Figure 1 and the topology characteristics from Figure 2.

3.1. Network Paths and Hoplists

The building blocks of a topology graph are the paths between nodes. Each path identifies a source and destination node. Paths are unidirectional for topology purposes. A
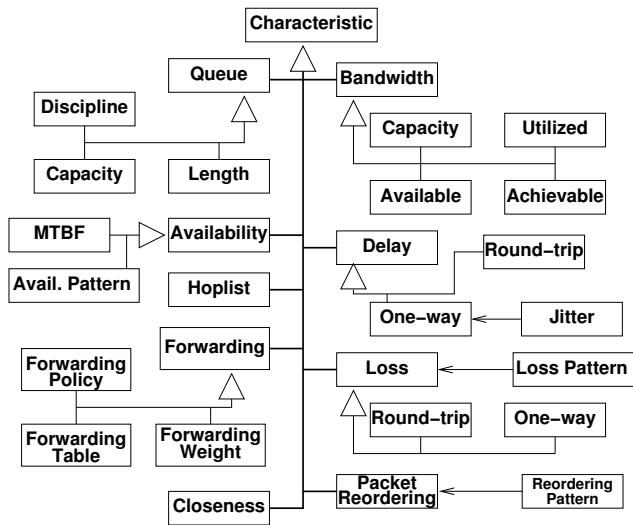
Figure 2: Network characteristics that can be used to describe the behavior of network entities.

path may be either a host-to-host path or a single "hop." A measurement system reporting only characteristics of end-to-end measurements would use a single host-to-host path for each observation, without relating the multiple paths. On the other hand, a system representing an entire topology could use router-to-router and router-to-host paths as the edges forming the topology graph connecting all of the hosts and routers.

In many situations, knowledge of the components that make up a path is useful. On host-to-host paths, meaningful measurements can be taken while only the endpoints are known. However, observing the hops between hosts, e.g. as reported by traceroute is frequently useful. The *hoplist* characteristic, shown in Figure 2, provides this information. For example, a path between two end-hosts could have a hoplist containing the set of router-to-router hops discovered by traceroute. Each member of the hoplist is itself a *hop*, which inherits from *path* as shown in Figure 1. End-to-end paths are not the only paths that can be divided into hops with a hoplist—one of the above router-to-router hops could be further subdivided into hops representing the links between Ethernet switches in a LAN.

The options of describing either complete topology graphs or only host-to-host connections, with or without hops specified, provide a powerful set of representations covering the various topology representations used by current network measurement systems.

3.2. Physical and Functional Topologies

Monitoring systems use different techniques to discover and represent the network between end-hosts. We provide

physical nodes, such as routers and switches, to represent physical topology and virtual nodes to represent functional or approximate topology.

The physical approach determines the actual physical links, including switches that are the internal nodes along the path. The physical approach, when combined with details about capacity and load, allows detailed modeling of LANs [2, 14] and WANs [8, 18].

The functional approach makes use of end-to-end information. Functional topologies can be used to represent connections according to perceived performance or to represent shared-media networks that are not amenable to a typical graph representation. The functional approach uses virtual nodes to build a topology and can be built between a variety of sites distributed around the Internet or using from the perspective of a single measurement node [3, 11, 21, 24].

4. Performance Characteristics

The following section briefly describes several important network performance characteristics. More detailed descriptions are available in [15].

Bandwidth Bandwidth is defined most generally as data per unit time. However, the "bandwidth of a path" is not a precisely defined term, and specific characteristics must be clearly defined prior to discussing how to measure bandwidth. There are four characteristics that describe bandwidth: capacity, utilization, available bandwidth, and achievable bandwidth.

Capacity is the maximum amount of data per time-unit that the path has available when there is no competing traffic. The capacity of the bottleneck hop is the upper bound on the capacity of a path that includes that hop.

Utilization is the aggregate capacity currently being consumed on a path over a particular time interval. Available bandwidth is the maximum amount of data per time-unit that a path can provide, given the current utilization.

Achievable bandwidth is the maximum amount of data per time-unit that a path can provide to an application, given the current available bandwidth, the protocol, the operating system used, and the end-host performance capability and load. While the bottleneck hop may limit the achievable bandwidth, hardware configuration or software load on the end hosts can possibly (and even likely on high-speed networks) limit the achievable bandwidth delivered to the application.

Delay and Jitter Delay is the amount of time for a packet to travel from the source to the destination. Delay can be measured one-way or round-trip. One-way delays are important because today's Internet connections often use

4

asymmetric paths, which are often different lengths. Even symmetric paths may have different characteristics due to asymmetric queuing. For more information see RFC2679 and RFC2681.

The variation in the delay of packets is sometimes called the "jitter." Because it is a partial statistical description of another characteristic, we place it under delay in the characteristics hierarchy.

Loss Loss is a measure of the number of packets that are sent out but never make it to the destination. The percentage of lost packets is the simplest measurement, but loss patterns, such as loss burstiness, are also important to protocols' performance. For more information see RFC2680 and RFC3357.

Availability The network availability characteristic is a measure of how often the network is "up" or "down." The exact definition of "down" compared to "very congested" may vary and probably depends on the application. For example, one definition may be that if no ICMP packets get through the network for 10 seconds, the network is considered "down".

5. Evaluation

We will evaluate the network measurement hierarchy by addressing the following questions:

- Do existing tools map to the hierarchy?

- Is the hierarchy useful to existing monitoring systems?

- Does the hierarchy simplify comparing the results from existing tools?

First, we looked at 19 currently used network measurement tools and verified that the results from each of these tools could be described using one of our characteristics [20]. Next, we looked at several existing monitoring systems, including NTAF, IEPM, GridLab, and NWS. These systems have all been modified to use the new NMWG characteristics hierarchy for naming the results. These systems are described in more detail below, along with a description of how the NMWG characteristics are used. Finally, we investigated issues of result compatibility between tools measuring the same or related characteristics.

5.1. Measurement Tools

In terms of Figure 1, the 19 tools implement measurement methodologies that in turn measure certain network characteristics. We were interested in checking whether the results of existing tools can be expressed using our hierarchy or whether mismatches or omissions exist.

The tools investigated cover the four bandwidth characteristics as well as round-trip delay, loss, and hoplist. Some tools report results for complete end-to-end paths, others report characteristics on a hop-by-hop basis. In total, nine different measurement methodologies are implemented, having varying levels of precision and intrusiveness.

We can summarize our study by observing that our hierarchy can indeed express the characteristics measured by the tools we investigated. No significant mismatch or omission has been found. Instead, our hierarchy models many additional potentially interesting characteristics.

5.2. Implementations

We now describe several monitoring projects that are using or starting to use these network characteristics. The goal of these systems is to make it easy to collect, publish, and compare results from any set of network or host monitoring tools running at multiple sites in a WAN. The basic function performed by the systems is to run network monitoring tools at regular intervals and store their results in an archive system (either centralized or distributed) for later analysis. Recent data are cached and published using the names and attributes being defined by the GGF NMWG profiles document [17].

NTAF One system currently using the hierarchy is the Network Tool Analysis Framework (NTAF) developed at Lawrence Berkeley National Lab. NTAF provides full mesh network testing using a configurable set of measurement tools.

The NTAF system has modified or wrapped the following network test tools to generate monitoring events using the standard names and attributes described in this paper: ping, iperf, netest, pathrate, and pathload. NTAF is also running a GridFTP client to test disk-to-disk achievable bandwidth. NTAF servers have been deployed at several sites. The network characteristics currently being collected and published, along with the tools currently being used to collect this information, are shown in Table 1.

All NTAF-generated monitoring events contain the following information: timestamp, program name, event name, source host, destination host, and value. For extensibility, arbitrary additional data can be included with each event. Using a standard event format with common attributes for all monitoring events allows us to quickly and easily build SQL tables of the results. More details are in [12].

IEPM-BW and PingER Another similar project is the Internet End-to-end Performance Monitoring—BandWidth

Table 1: Characteristics and tools used by different projects

NTAF

`path.delay.roundTrip`	ping
`path.delay.oneWay`	netest
`path.loss.oneWay`	netest
`path.bandwidth.achievable.TCP`	iperf, netest, GridFTP
`path.bandwidth.achievable.TCP` `.multiStream`	iperf, netest, GridFTP
`path.bandwidth.achievable.UDP`	netest
`path.bandwidth.available`	pathload, netest
`path.bandwidth.capacity`	pathrate

IEPM-BW and PingER

`path.delay.roundTrip`	ping
`path.delay.oneWay`	RIPE-tt
`path.delay.jitter.roundTrip`	ping
`path.loss.roundTrip`	ping
`path.availability.roundtrip`	ping
`path.bandwidth.achievable.TCP`	iperf, bbftp, GridFTP
`path.bandwidth.achievable.TCP` `.multiStream`	iperf, bbftp, GridFTP
`path.bandwidth.available`	ABwE
`hop.bandwidth.capacity`	SNMP
`hop.bandwidth.utilized`	SNMP

GridLab

`path.delay.roundTrip`	NWS
`path.bandwidth.achievable.TCP`	NWS
`path.bandwidth.capacity`	pathrate
`path.hoplist`	TopoMon

Figure 3: Example for bandwidth measured by different tools, on the path from LBNL to PSC

GridLab The European Union project GridLab [1] is monitoring the network between its testbed sites and reporting the measurement data based on the characteristics hierarchy presented here. Within the GridLab testbed, minimizing the intrusiveness of the measurements is important, so measurement data are stored where they are produced—in a distributed manner at each testbed site. A centralized web service lets application programs and the GridLab web portal retrieve the data on demand. The web service is soon to be replaced by an OGSA service. Minimizing intrusiveness also directed the selection of the NWS network sensor to measure `path.bandwidth.achievable` and `path.delay.roundTrip`, as well as the choice of pathrate and TopoMon [7], as shown in Table 1.

Network Weather Service The Network Weather Service [26] provides a caching proxy component called the NWSlapd [23]. This component supports the delivery of forecasts using a slightly older version of the measurement hierarchy described here. Once the standard is finalized, the NWSlapd will deliver forecast data in compliance with it.

5.3. Compatibility

Figure 3 shows an example of compatibility problems between various network measurement tools, as well as how the network characteristics hierarchy facilitates analysis and comparison among their results. The graph shows a comparison of `path.bandwidth.available` results for pathload and netest, and `path.bandwidth.achievable` for TCP and UDP, as measured by iperf and netest. All of the tools are run periodically over a 24 hour interval, and all TCP streams are tuned to use optimal buffer sizes.

For this particular path, LBNL to PSC (bottleneck hop = 600 Mbps), pathload gives a very consistent result of

(IEPM-BW) and PingER systems developed at SLAC. PingER makes ping measurements at 30 minute intervals from over 35 sites in 15 countries to over 400 sites in more than 75 countries. The data are collected by two archive sites at SLAC and FNAL.

IEPM-BW makes ping, traceroute, iperf, bbftp and GridFTP measurements, at 60 or 90 minute intervals, from 10 sites in 5 countries to about 50 sites in 9 countries. These measurements are stored locally on the measurement hosts. Table 1 shows the network characteristics being collected and published, along with the tools that are used.

Note that several tools generate the same information. For example, both pathload and netest measure `path.bandwidth.available`. The use of a common measurement hierarchy and naming system allows us to easily compare the results of these two tools, compare their results with `path.bandwidth.achievable` data from iperf, and make conclusions about which tool is more suitable for a given network path.
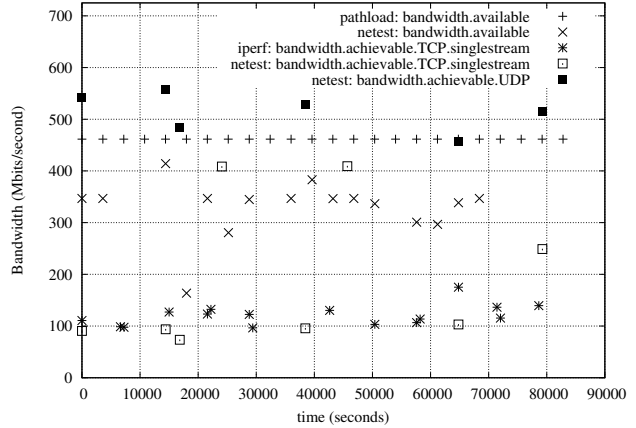
460 Mbps for available bandwidth, while netest gives values that vary between 150 and 400 Mbps. Not surprisingly, single stream TCP achievable bandwidth as measured by iperf is much less, around 100 Mbps for a single stream, and 150–250 Mbps for four parallel streams. netest results are usually similar, but occasionally seem to be too high. netest also provides a UDP achievable bandwidth measurement of around 400–500 Mbps.

Analysis of these results is far beyond the scope of this paper, but clearly the ability to assign standard names and characteristics to the results from these tools facilitates this kind of analysis. More generally, the following kinds of compatibility problems arise.

Same Characteristic, Different Methodology/Tool For a single characteristic, different measurement methodologies may produce different results. For example, `path.delay.roundTrip` may be measured using ping, using the smoothed RTT reported by Web100, or using a tool such as synack to measure the delay between a TCP SYN and the SYN/ACK returned. Differences in the results are presumably due to the ways in which ping and TCP are implemented, as was observed in [6].

Derived Characteristics For some characteristics, it may be possible to derive them from related characteristics rather than measuring them directly. For example, `path.bandwidth.achievable.TCP` may be derived from the loss and RTT measured by ping. This estimation can then be compared with the achievable TCP bandwidth measured by iperf. For paths with low losses, this approach may require considerable time to measure the loss rate accurately, and the ping loss may still not accurately reflect the TCP loss (e.g., since standard TCP deliberately provokes loss).

Comparing Different Characteristics As shown in Figure 3, care has to be exercised when comparing different characteristics such as `path.bandwidth.achievable.TCP` measured by iperf and `path.bandwidth.available` measured by pathload. Not only are the tools operating at different network layers, but also TCP transfer rates can be very dependent on configuration parameters, such as the maximum TCP window size and number of streams, whose optimum values vary with time (see for example [5]). However, the characteristics hierarchy enables such comparisons by clearly identifying the different things being measured.

6. Related and Future Work

The nomenclature and hierarchy presented in this paper provide a common dictionary of terms as well as relationships between commonly used measurements. The hierar-

chy allows measurements to be grouped according to the network characteristic they are measuring.

The NMWG focuses on existing and currently used measurements to determine network performance. It does not attempt to define new standards or to define only the best measurement methodologies to use for grid applications; however, it does attempt to point out the advantages and disadvantages of different measurement methodologies [15]. For example, the results from some measurement tools are influenced by bottlenecks at the hosts making the measurements. Identifying and resolving such bottlenecks is the focus of the Internet2 End-to-end Performance WG [10] and of part of the EU DataGrid and DataTAG projects [9].

The NMWG is closely related to the IETF Internet Protocol Performance Metrics (IPPM) WG. Whereas their work focuses on the best-practices for network engineers and defining the *right way* to take measurements, our goal is to create a comprehensive system that can be used to categorize all measurements that *are* in use, taking into account the requirements of grid applications and grid-related tools. Where possible, we adopt the terminology defined in the IPPM Framework[19], although due to the different goals of NMWG and IPPM, certain aspects of that framework do not apply to this document, such as the flat namespace and use of the term "metric." In practice, we find that the terms "characteristic," "measurement methodology," and "observation" are rarely misinterpreted, therefore, we prefer that terminology and avoid the difficulty of conflicting uses of the word "metric." We preserve the use of the other terminology described by RFC2330 wherever possible.

The authors have begun working on a common SOAP-based API for their measurement systems. The combination of the standard measurement characteristics hierarchy described in the paper, along with a standard schema for describing the attributes for each characteristic, will allow us to build such a standard API. The NMWG is currently defining this schema [17].

7. Conclusions

Measurement and prediction of network resources are crucial so that adaptive applications can make use of Grid environments. For grid middleware to make sense out of the results from a wide variety of network measurement systems, a standard for describing network measurements is required. In this paper we have summarized our work within the Global Grid Forum's Network Measurements Working Group, which has led to defining a standard for representing network entities and measurements of their properties. This standard enables the exchange of network measurement data between multiple Grid Services. We have also demonstrated that this measurement hierarchy is useful for

storing and exchanging measurement data between various systems.

Acknowledgements

References

[1] G. Allen, K. Davis, K. N. Dolkas, N. D. Doulamis, T. Goodale, T. Kielmann, A. Merzky, J. Nabrzyski, J. Pukacki, T. Radke, M. Russell, E. Seidel, J. Shalf, and I. Taylor. Enabling Applications on the Grid: A GridLab Overview. *International Journal of High Performance Computing Applications*, 2003.

[2] Y. Breitbart, M. Garofalakis, C. Martin, R. Rastogi, S. Seshadri, and A. Silberschatz. Topology discovery in heterogeneous IP networks. In *INFOCOM 2000*, March 2000.

[3] M. Coates, A. O. H. III, R. Nowak, and B. Yu. Internet tomography. *IEEE Signal Processing Magazine*, 19(3):47–65, May 2002.

[4] R. L. Cottrell and C. Logg. A new high performance network and application monitoring infrastructure. Technical Report SLAC-PUB-9202, SLAC, 2002.

[5] R. L. Cottrell and C. Logg. Experiences and results from a new high performance network and application monitoring toolkit. In *PAM2003*, 2003. SLAC-PUB-9560.

[6] R. L. Cottrell and M. Shah. Measuring rtt by using syn/acks instead of pings. http://www-iepm.slac.stanford.edu/monitoring/limit/limiting.html#synack, December 1999.

[7] M. den Burger, T. Kielmann, and H. E. Bal. TopoMon: A monitoring tool for grid network topology. In *International Conference on Computational Science (2)*, pages 558–567, 2002.

[8] R. Govindan and H. Tangmunarunkit. Heuristics for internet map discovery. In *IEEE INFOCOM 2000*, Tel Aviv, Israel, March 2000.

[9] R. Hughes-Jones, S. Dallison, G. Fairey, P. Clarke, and I. Bridge. Performance Measurements on Gigabit Ethernet NICs and Server Quality Motherboards. In *PFLDnet 2003*, 2003. http://datatag.web.cern.ch/datatag/pfldnet2003/.

[10] Internet2 end-to-end performance initiative. http://www.internet2.edu/e2eperf/.

[11] S. Jamin, C. Jin, Y. Jin, D. Raz, Y. Shavitt, and L. Zhang. On the placement of internet instrumentation. In *IEEE INFOCOM 2000*, Tel Aviv, Israel, March 2000.

[12] J. Lee, D. Gunter, M. Stoufer, and B. Tierney. Monitoring data archives for grid environments. In *IEEE Supercomputing 2002 Conference*, 2002.

[13] B. B. Lowekamp, N. Miller, R. Karrer, T. Gross, and P. Steenkiste. Design, implementation, and evaluation of the remos network monitoring system. *Journal of Grid Computing*, 1(1), 2003.

[14] B. B. Lowekamp, D. R. O'Hallaron, and T. Gross. Topology discovery for large ethernet networks. In *SIGCOMM 2001*. ACM, August 2001.

[15] B. B. Lowekamp, B. Tierney, R. L. Cottrell, R. Hughes-Jones, T. Kielmann, and M. Swany. A hierarchy of network measurements for grid applications and services. Draft-GGF-NMWG-Hierarchy, 2003. http://forge.gridforum.org/projects/nm-wg/.

[16] W. Matthews and R. L. Cottrell. The pinger project: Active internet performance monitoring. *IEEE Communications Magazine*, pages 130–137, May 2000.

[17] Network Measurements Working Group. Schema/profile for network performance measurements for grids. `http://www-didc.lbl.gov/NMWG/NMWG-profile.html`.

[18] V. N. Padmanabhan and L. Subramanian. An investigation of geographic mapping techniques for internet hosts. In *ACM SIGCOMM 2001*, pages 173–185, 2001.

[19] V. Paxson, G. Almes, J. Mahdavi, and M. Mathis. Framework for IP performance metrics. RFC2330, May 1998.

[20] R. Schoonderwoerd. *Network Performance Measurement Tools - a Comprehensive Comparison*. M.Sc. Thesis, Vrije Universiteit, Amsterdam, November 2002. available from `http://www.cs.vu.nl/albatross`, summary results available at: `http://www-didc.lbl.gov/NMWG/NM-WG-tools.html`.

[21] G. Shao, F. Berman, and R. Wolski. Using effective network views to promote distributed application performance. In *International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA'99)*, 1999.

[22] M. Swany and R. Wolski. Multivariate resource performance forecasting in the network weather service. In *SC02*, Baltimore, MD, November 2002.

[23] M. Swany and R. Wolski. Representing dynamic performance information in grid environments with the network weather service. 2nd IEEE International Symposium on Cluster Computing and the Grid, May 2002.

[24] W. Theilmann and K. Rothermel. Dynamic distance maps of the internet. In *IEEE INFOCOM 2000*, Tel Aviv, Israel, March 2000.

[25] B. Tierney, W. Johnston, B. Crowley, G. Hoo, C. Brooks, and D. Gunter. The netlogger methodology for high performance distributed systems performance analysis. In *Seventh IEEE International Symposium on High Performance Distributed Computing (HPDC 7)*, pages 260–267, 1998.

[26] R. Wolski. Dynamically forecasting network performance using the network weather service. *Cluster Computing*, 1:119–132, January 1998.