SLAC-PUB-95-7008
November 1995

DISTRIBUTED COMPUTING ENVIRONMENT MONITORING AND USER EXPECTATIONS*

R. L. A. COTTRELL, C. A. LOGG

Stanford Linear Accelerator Center, Stanford University, Stanford, CA 94309, USA.

This paper discusses the growing needs for *distributed system* monitoring and compares it to current practices. It then goes on to identify the components of *distributed system* monitoring and shows how they are implemented and successfully used at one site today to address the Local Area Network (LAN), network services and applications, the Wide Area Network (WAN), and host monitoring. It shows how this monitoring can be used to develop realistic service level expectations and also identifies the costs. Finally, the paper briefly discusses the future challenges in network monitoring.

1 Introduction

The HEP community is rapidly moving to a client/server distributed environment. Though there are many benefits in such an environment, there are major challenges to deliver the levels of *distributed system* service that users expect and will need in the increasingly global environment of the next millennium.

These user expectations are often based on the local performance of a stand alone workstation or a mainfarme. These have a single operating system, file system, set of system services and physical backplane. In comparison, today's networked environment is very complex.[1] It includes equipment and software from multiple vendors, multiple protocols, distributed file systems and system services, and a "backplane" with both hardware and microcode components which is constantly being modified by people with widely varying skills and responsibilities. This complexity has resulted in decreased support effectiveness and buying power, which in turn leads to: decreased quality of service; inability to support existing and new applications; and increased downtime, users' time wasted and security exposures.

Subjectively, the ultimate measures of *distributed system* performance are the users' perceptions of the performance of networked applications. Examples of such applications are WWW, Email, a distributed database, or a spreadsheet accessing a distributed file system. *Distributed system* performance is affected by the physical network plant, computers and peripheral devices attached to it, and the software (from device interface through application) running on the computers and devices. To set and meet user expectations for system performance we must monitor all the above components.

At the same time, budgets are increasingly constrained, 55% of companies indicate[2] they are understaffed for managing network performance. There is a severe lack of experienced personnel. It is also hard to retain trained personnel. Networks are growing in extent, numbers of devices (30-50% growth/year is typical), traffic (typically doubling

*. Work supported by Department of Energy contract DE-AC03-76SF00515.

Contributed to the International Conference on Computing in High Energy Physics (CHEP95), Rio de Janeiro, Brazil, September 18-22, 1995.

every 18 months), and the technology to manage the network is not growing as fast as the network technology itself. With these challenges it is critical to provide easy-to-use, integrated tools, to automate the monitoring and enable effective management of the heterogeneous *distributed system* environment with the limited resources available. McConnell Consulting, for example, estimates that by using RMON[4] tools organizations can double the number of nodes a management staffer can manage.

2 Why Monitor

The top reason (10 on a scale of 1-10, with 10 the most important) for monitoring the network today, according to a May-95 survey of 9 ESnet sites[5]representing a total of over 50,000 networked devices, was unanimously agreed to be to assist in **trouble shooting**. This is needed to get out of crisis mode, to identify problems and start diagnosis/fixing problems before the users notice, and to allow the users to accomplish work more effectively and maximize their productivity. This was closely followed by **performance tuning** (8.4) which allows networkers to proactively reduce bottlenecks, tune and optimize systems, improve the quality of service, identify under and over utilized resources and balance work loads. This in turn was followed in close succession by **planning** (8.0) to understand performance trends, setting and tracking **expectations** (7.3), and security (7.2). **Accounting** brought up the rear (4.0). According to a recent survey by McConnell Consulting, we may expect accounting to almost double in importance in the next 2 years, so it will be more in-line with the other reasons to monitor.

3 What Should We Monitor and Current Practices

The ESnet survey, mentioned above, also provides input on what we actually monitor and what we want to monitor. The results are seen in Figure 1. All sites monitor the status of network devices (up, down etc.), about 50% monitor the status of the links (utilization,

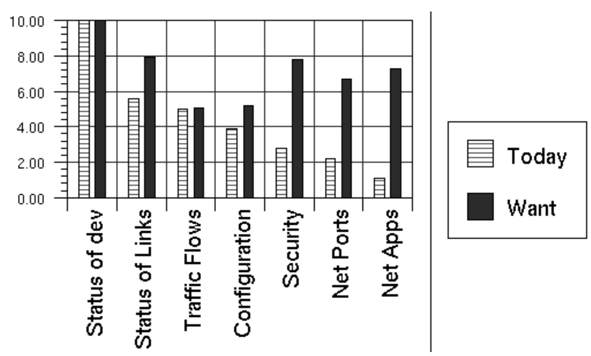

Figure 1: What ESnet Sites Monitor Today vs. What They Want to Monitor

errors, etc.) and the traffic flows (network level accounting, who talks to whom). Less than 40% of the sites actively monitor the configurations (who is where and at what address), and even fewer monitor security and network applications (either the ports/daemons or the applications themselves). The greatest desires appear to be to increase application and security monitoring.

The difficulty that companies encounter in trying to manage network performance is also illustrated in a recent survey[1] by INS. This survey indicates that companies report that: only 24% adequately manage network performance; 95% would like to report on network utilization, but only 55% do; 91% would like to report on network availability, but only 25% do; 56% have a project in the works or plan to improve network performance; 65% have a project in the works or planned to improve network management; and only 16% have network performance service level agreements.

4 How Network Monitoring is Accomplished Today

The many components of network monitoring and their inter-relationships are illustrated in Figure 3. We will next show how these components are used with particular reference to current practices at SLAC[3].

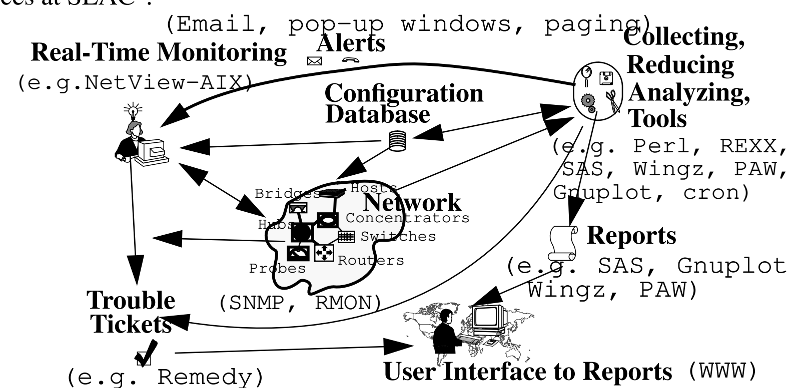

Figure 2: The Components of Network Monitoring and their Inter-relationships

4.1 Data Collection, Analysis and Reduction

Data are collected at regular intervals from the Simple Network Management Protocol (SNMP) agents' Management Information Bases (MIB) in the routers, bridges, probes, hubs and switches. The data collected include: # good packets; # kilobytes; packet size; # errors; # packets dropped, discarded, buffer controller overflows; top-10 talkers and protocol distributions; and address information from the Address Resolution Protocol (ARP) caches. Data are also collected using the IP `ping` function to measure the response time, packet loss and connectivity of critical servers, routers interfaces, probes and off-site collaborators' nodes. Finally we poll critical network server daemons and services such as name, font, file service, WWW and Email.

The data collected are analyzed by batch jobs at weekly, daily and hourly intervals, based on the time scales needed for the reports. The analysis summarizes the raw data into ASCII files (usually tabular reports) and postscript graphs showing daily, fortnightly, monthly and longer term trends.

The analyses generate thousands of reports, most of which are uninteresting. The next step is therefore to automatically examine the ASCII tabular reports and extract the exceptions. These exceptions include factors like: duplicate IP addresses from the ARP caches; the appearance of unregistered nodes; loss of connectivity discovered by `ping` or SNMP polling; data values exceeding thresholds such as frame checksum and alignment errors > 1 in 10,000 packets, total utilization on a subnet of > 10% for the day, a broadcast rate of > 150 packets/sec, (shorts+collisions)/GoodPackets > 10%, packet loss from pings to on site nodes of > 1% in a day; bridge and router overflows and queue drops, and poor response from network services.

This examination of exceptions is used to generate exception reports for display by WWW. The reports display the exceptions as hypertext links to tables and graphs with more information.

4.2 Alert Notification

The daily WWW visible exception reports are manually reviewed each workday morning and used as input to the Happening Out There (HOT) meeting. The HOT meeting is a 5-15 minute open meeting of responsible people from network operations and development, systems administration, the help desk and interested users. It covers scheduled outages, installations, newly encountered problems and outstanding/unresolved problems.

Critical alerts from the SNMP and `ping` polling of critical interfaces also results in issuing X-window pop-up windows on selected screens, phone pages being issued and Email messages being sent to the relevant people.

Real-time alerts are also provided by a DECmsu network monitoring platform which displays maps of the Local Area Network (LAN) and displays nodes in red when they are unreachable. This map is monitored by the help desk during normal working hours. The map is also available to critical support people at their workstations.

4.3 Wide Area Network (WAN) Monitoring

WAN monitoring for an end site has different requirements than LAN monitoring since the management of the WAN is usually outsourced to a network provider, so the end site has little control. In addition there is much greater variability in measurements such as response time. However, users and networkers still want to have reasonable expectations for performance, planning and problem identification.

The main tools used today by end sites such as SLAC are: `ping` to determine response time, packet loss, and connectivity; `traceroute` to discover the routes between nodes; measurements of file transfer rates; and traffic measurements at the firewall router. Several HEP sites such as IN2P3, Padua, RAL and SLAC have automated the pinging and produce reports[†].

Using these tools we: look for significant changes in response time which may indicate routing problems; separate packet loss (intermittent) from node/link down (no

†. See for example http://www.slac.stanford.edu/comp/net/wan-mon.html.

response for long periods); and, look for significant increases in packet loss that may indicate overloaded routers/links. Understanding the file transfer rates is more complex since it involves many factors including packet loss, network response time, number of hops and route used, utilization of links, speed of links and the capability of the end nodes. For example, we see typical changes in average transfer rates of 30% between holidays and work days, and even changes of a factor of 5 or more from minute to minute for a given node.

Since the quantitative values of the measurements (response, packet loss, transfer rates) differ widely from node to node and even month to month, we are working on developing automatic methods to set dynamic thresholds for alert generation.

4.4 Host Monitoring

It is becoming increasingly important to also monitor the resources of the host systems and mission-critical applications components of the *distributed system*. Though there are a number of relevant experimental and standards-track MIBs including the Host Resources MIB[6], the developers of device drivers, operating systems and system daemons must provide means of accessing the information. Thus much of today's monitoring practice is roll-your own and system specific. At SLAC we perform extensive accounting/monitoring on our Unix servers to provide reports on: disk space, memory and cpu utilization; paging activity and input/output; account utilization; and security intrusions such as a host's Ethernet interface being discovered in "promiscuous" mode, or if the enciphered checksums of critical system files have changed. These reports are accessible via WWW, and automatic alarms are generated, for example on-call people are paged in case of a security intrusion, and users are warned by Email if their application on a server uses > 10% of the cpu for > 15 minutes.

4.5 Expectations and Costs

Using the results of the monitoring we have been able to set baselines, and publish service level expectations. Examples of the service level expectations today are: `ping` response for the on-site network layers < 0.01 seconds for 95% of the samples; network reachability of >= 95%; sub-second response for trivial network services (name, font, network daemons such as smtp, nfsrpc); 95% of trivial mail deliverd on site in 10 minutes; 95% of requests for the SLAC WWW home page served in < 0.1 seconds.

Polling of devices can utilize significant bandwidth. For example SNMP polling 20 stations/interfaces/agents every 5 seconds can use all of a 64 Kbps link[‡]. Care must therefore be taken to: select only the essential variables; match the polling volume to node/path, ensuring the polling is minimized for remote sites; and to consider the time relevance of the data, for example accounting information capability is unlikely to be needed more frequently than a daily basis.

‡. SLAC with about 1000 interfaces which are polled, uses < 0.3% of the Ethernet bandwidth on average.

In terms of disk space, SLAC gathers about 4 megabytes/day of raw data, and has about 750 megabytes of graphs on-line. There are about 16,000 lines of code mainly in SAS, Perl and REXX. Of this 77% is for the analysis, 15% for the collection and 8% for the reduction phases. Parts of the code has been exported to other sites, and there has been considerable interest from other sites and vendors. The hard part is not in the coding, but in defining and understanding the details. In all the effort has been about 3 full time equivalent people-years.

5 Conclusion

There is still no out-of-the-box integrated solution available for network monitoring. The information provided today needs to be made much more digestible. Network managers need to encourage vendors to provide the tools, for example by demanding SNMP and RMON agents in all appropriate devices.

Currently we are extending the SLAC system to cover the Fiber Distributed Data Interface (FDDI) backbone and the newer Ethernet and FDDI switches. For the future we look forward to the provision of better standards-based monitoring tools (e.g. based on RMON 2[**]) for the higher layers of the OSI network stack. Beyond that the provision of monitoring for Asynchronous Transfer Mode (ATM) networks promises to be a major challenge which is only just starting to be addressed in the standards bodies.

Developing tools to effectively monitor is costly and still has to be done in-house.Yet, we strongly believe that good monitoring tools can effectively leverage scarce resources and enable setting of realistic service level expectations. WWW has been an outstanding vehicle for providing the user interface, though there are concerns about the granularity with which one can restrict the visibility of some sensitive reports.

Acknowledgments

We would like to thank our many colleagues at SLAC, the ESnet community and vendors who have contributed to the success of this system. In particular we acknowledge the major contributions of Lois White, Mike Sherman, John Halperin, Teresa Downey, Don Pelton and Chuck Boeheim.

References

1. Nick Lippis, John Morency and Eric Hindlin in *ConneXions,* Vol 5, 18 (1995).
2. Jeff Paschke *Network Performance Survey Results Report*, from International Network Services, (June 1995).

**. The Internet Engineering Task Force published RMON version 2 as a proposed standard and RFC in July 1995.

3. C. A. Logg and R. L. A. Cottrell *Network Management and Performance Monitoring at SLAC*, SLAC-PUB-95-6744 and Presented at Networld + Interop Conference, Las Vegas, (March 1995).
4. Steve Waldbusser, *Remote network monitoring management information base*, RFC 1271, (Nov 1991).
5. R. L. A. Cottrell et. al. see `http://www.slac.stanford.edu/~cottrell/tcom/nmtf.html`.