

Evolution of a Network Management System

Anthony Michel

BBN Systems and Technologies

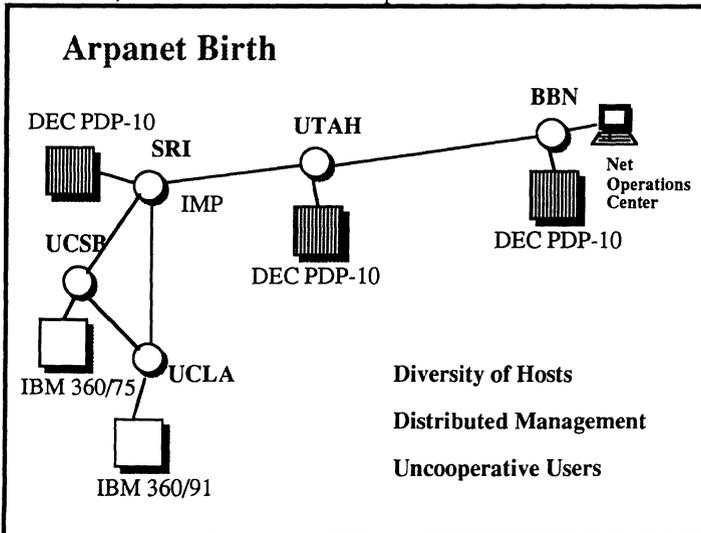
Cambridge MA USA michel@bbn.com

Abstract

Packet net management architecture and technology have evolved out of the earliest work on the Arpanet. Net management techniques developed for the Arpanet have become the foundation for many modern NM systems. At present, this field is intensely competitive and in the future participants will require attractive technology to succeed. Some of the technology that will be important in future systems, such as CORBA for information distribution, is being developed now, and this will help to make use of the great amount of existing technology that has been developed in the Internet community. Some areas need more fundamental research, for example on generating leading indicators for network trends and. Additional work is needed on efficient ways to collect data on the "vital signs" of a network - which is like a living organism.

1. Arpanet History

25 years ago the first packet switch, called an IMP[1], started service at UCLA. IMP1 was soon joined by 4 more nodes, which became the initial Arpanet.

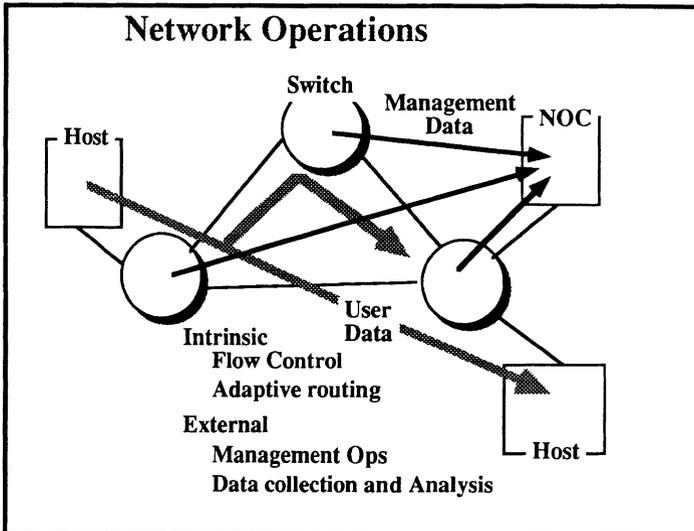


The Arpanet was a big success and this technology has had a powerful influence on many subsequent data communications systems. Among the reasons for this success, the original developers single out the "remote net management" feature of the Arpanet. This capability for completely autonomous, self-configuring operation was quite unusual in 1969, and it was necessary. The first network nodes were in California, while the development and operations center was in Massachusetts.

The Arpanet technology was based on several ideas that became practical as the cost of small, powerful computers declined in the late 1960's. Working in Japan, Hasegawa [2] and his group advocated sharing a communications link among a group of contending users, by sending "blocks" of data from various users on a single link, and allocating the link only for the duration of the block. In the USA, Baran [3] advocated the use of small computers to execute the nodal switching functions and manage the mis-match of data rates inherent in such systems. The US Department of Defense Advanced Research Project Agency (ARPA), launched a series of development efforts to build a prototype network based on the idea of "Packet Switching" and to develop a family of host computer protocols and applications which could efficiently use the net.

1.1 Arpanet Elements

The net management system for the Arpanet, called the NOC, or Network Operations Center[4], was responsible for operational supervision and troubleshooting of the switches and communications links. In addition, it helped developers at user sites who were developing and debugging network software. Three basic elements comprised the original Arpanet: hosts, IMPs, and links.



Hosts:

Computers, called hosts, were the users of Arpanet communications. The first 2 hosts were IBM 360 mainframes, and these were soon joined by DEC PDP-10s. The DEC and IBM architectures were very different, and they shared no useable communications protocols. It was necessary to develop an entirely new communications architecture, which was called the "host to host protocol."

Switches:

The packet switches of the Arpanet, called IMPs, were built on Honeywell 516 minicomputers. The software was a completely new design, done by Bolt Beranek and Newman (BBN). The switches were designed to support full-duplex transmission links at a data rate of 50,000 bits/sec, at a time when computer communications was mostly done at 2400 bits/sec, half-duplex. A whole new switch architecture was needed.

Communications Links:

The interswitch trunk lines and modems were provided by AT&T Long Lines. Full duplex leased circuits of 50Kilobits/sec, a very high data rate at that time, were employed. The circuits were the most expensive part of the network, and so it was very important to keep them in service at all times. Although it was feared that these links would be error-prone and troublesome, this was not the case. The bit error rate of the lines was unmeasurably low, except for the "bad lines." Bad lines were easily identified by their exceedingly high error rate, which could usually be traced to faulty power supply equipment.

1.2 Arpanet Overview

There were two purposes in constructing the Arpanet. First, was to test the idea of packet switching on a large scale with real users and circuits. The second idea arose out of the first, and it was to provide a computer resource sharing environment for researchers, to save money. The idea was to collect computing resources into several large centers, and then give access to these centers to researchers all across the country. The packet switching test led us to build considerable statistics collection machinery into the IMPs and the desire for reliable operation forced us to build in self-test and support functions. The IMPs were placed at user sites, sometimes in a computer room, but nearly always without much human operator supervision. The IMPs had to survive on their own. The IMPs had no direct human users, and a measure of their success was the degree to which they were invisible. So long as their host computers received reliable network service, no one at the site paid them much attention. From the beginning, very high reliability was achieved and users came to expect nearly 100% availability of the network.

1.3 Early Arpanet Management

The Arpanet network management system was a principal contributor to the good results. From the beginning, the NOC performed 3 main functions: troubleshooting coordination, software support, and network analysis.

- **Troubleshooting** - even though the IMP was designed for completely automatic, autonomous operation, hardware and software failure ultimately need human intervention. The IMP program included test programs built in to the main line operational S/W. Data consistency and program integrity checks were run continuously, even (especially!) under heavy traffic load. When measured values were found outside the expected range, the IMP sent "trouble reports" back to the NOC. This was immensely useful both for software and hardware problem detection. Usually we could notice a problem with a network node and contact the site before they were aware of the problem.
- **Software Support** - The IMP program was initially only 8000 words long - which comprised the entire software load for the machine. Nonetheless, it was quite complex and because it was highly optimized for speed (to minimize switch delay and maximize net throughput) it was difficult to understand and debug. We decided to concentrate all the programmer talent in one place, in the NOC at Cambridge, and give them tools to remotely observe and debug the IMPs, while they operated on the net. This was a most unusual project approach at the time, when a more customary approach was to put experts on site with complex computers, and deal with problems locally. With a central staff, we could afford to keep a few very highly qualified engineers dedicated to the problem-solving aspects of net operations.
- **Network Analysis** - There was considerable interest, from the start, in measuring all sorts of aspects of the Arpanet's performance. A center for theoretical analysis at UCLA focussed on numerical simulation and modelling of the Arpanet, while at BBN we concentrated on continuous collection of real, measured values for host throughput, link loading, end-to-end delays, and so

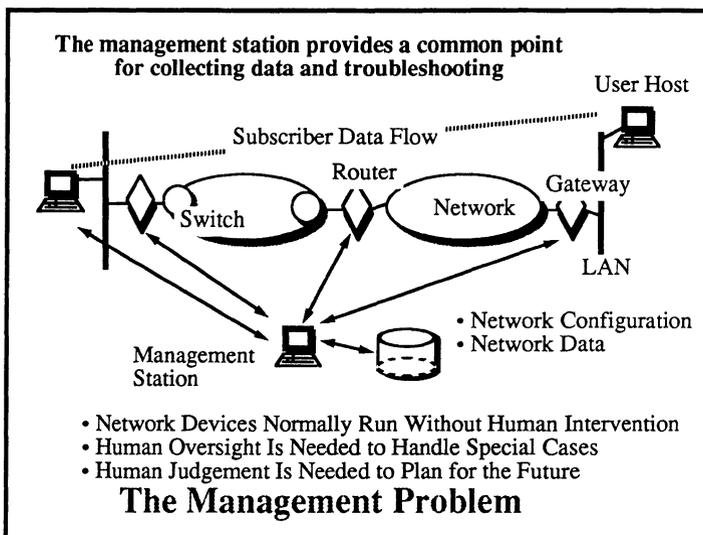
forth. These centers shared data and evolved results by collaborating over the Arpanet. In fact the first useful traffic on the net was transmission of host throughput data to the center at UCLA which owned Ursa, a large IBM 360/91 - then one of the fastest computers available.

1.4 Lessons from Arpanet Management

A principal lesson from the early Arpanet is that it is worthwhile including extra complexity in the network switches to support net management. The improvement in visibility into the switches, and the reduction of staff at remote sites, will pay for bigger switch computers.

2. Net Management Overview

BBN has about 100 network customers at present and they all have some interest in NM. Most customers are sophisticated and knowledgeable and usually each has a complex set of requirements. BBN is a supplier of network switching equipment, such as X.25 packet switches (successors to the IMP) and ATM switches and we must have an NM story for every customer. This is difficult, because the customers expect that all of their network elements will work together, and this means that many different NM architectures must work together. Committees from the ISO as well as the NM Forum have developed a general architecture for discussing NM. This is based, in part, on work done at BBN on status monitoring, analytical data collection, and troubleshooting and control functions of NM. For the most part, these functions are done by the humans, the network control operators who work in the NOC. NM is the function of the software that helps these humans.



3. Pressures from Customers

Now in the mid-1990's, the net management situation has reached a certain commercial and technical maturity. There are powerful products available from many reputable vendors. For a traditional vendor of network products and services, such as BBN the situation can be seen as:

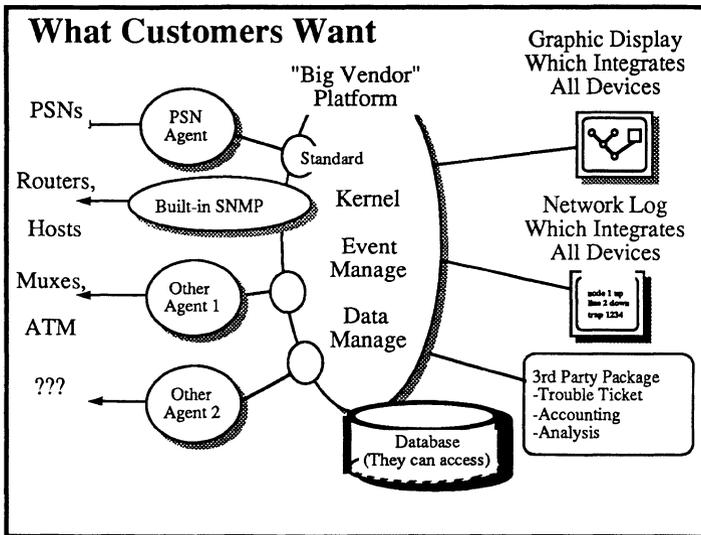
- Most organizations have a data communications network
- Internet architecture, based on TCP/IP is the dominant future information architecture
- Most organizations are building private NM systems
- If they obtain service from a public net, they still want the ability to observe and intervene

- Most organizations have opinions and plans for the structure of their NM system
- Customers are usually smart and well-informed about their options and constraints

3.1 What They're Asking For

Many potential customers seem to agree on a structure for the ideal NM system.

- It should work with one of the "Big Vendor" platforms, such as HP OpenView, IBM Netview/6000 or Sun Net Manager.
- It should be "open" and standards-oriented
- They wish to buy components from many suppliers to meet specific needs
- They expect to use the Internet architecture, with SNMP as the principal protocol
- It should support "Legacy" systems - older or existing systems that do not use SNMP or CMIP
- Very few networks outside of the telephone companies use CMIP
- They want to support a variety of analytical functions as well as element monitor and control
- They want to manage from several locations
- For reliability, they want to have several management centers
- Private network owners want to cooperate with the NM centers of their public net suppliers
- The central "WAN manager" must provide WAN visibility to "LAN managers."



3.2 Who They're Asking

To create the solution and build the system, customers have several options. In some cases they work with an NM product vendor such as HP or IBM; or they work with a vendor of the network devices, such as BBN. Traditional systems integrators and integrators who specialize in NM, as well as consultants also play a role. But the most important contributor in many organizations is often none of these. In many cases, the customers choose to develop their NM systems with their own staff. These systems are so important to an organization that their construction cannot be trusted to outsiders. In addition, NM systems are usually highly tailored to the needs of an organization. Each is a unique development, built to order for its owner.

4. BBN Approach

At BBN, we want to continue to develop, build and sell NM systems, and so we have looked hard

at where NM technology is going, and what potential customers want to buy. It is clear that the field is crowded and some of the players are large. One set of opportunities will be found in middleware, software which connects together the other pieces of the system. Another opportunity is supporting network analysis functions with data manipulation and statistical analysis functions. Additionally, BBN has always been strong in design and construction of special interfaces both hardware and software. This expertise is useful in building the specialized Device Agents. We plan to develop in these areas. We expect to work closely with big vendors like IBM, Sun and HP, each of whom have established positions as platform suppliers.

4.1 Overview of a Technical Approach

3 things influence our technical approach:

- Our legacy products

We have several thousand network devices fielded and these must be supported as the networks evolve with new devices and new NM systems.

- Our present expertise

We're especially good at certain things like network analysis. We hope to use this expertise to build software that could be used by customers to analyze their networks and understand the trends for traffic loads, network performance and user satisfaction.

- Where we are heading

BBN has several new and emerging products such as an ATM switch which require NM. We want to support our products and ensure that they are attractive to the widest possible market whether or not our net management platform is used. When we develop advanced analysis software or unique device agents, we want these to work with other NM platforms as well as our own.

4.2 Platform Approach

The structure of the platform we will use in the future resembles the architecture that we have been developing for many years. With this approach we can use device agents of various sorts.

- **Device Agents** - We have developed our own device agents for a number of BBN devices.

We can preserve these and augment them with several different SNMP devices. We can add device agents from elsewhere, to handle device types for which we have no support such as SONET multiplexors which report status with ASCII-coded alphanumeric event messages. In some cases, a general purpose platform like Sun Net Manager may be used as a device agent, if it exports management data to some other platform.

- **Presentation Agents** - Each user usually prefers a unique view of the data that NM generates.

The view that is suitable for a network schematic map of status is very different from a view suitable for tabular data. In many cases it may be desirable to use a general purpose NM platform like HP OpenView as the "viewer." An this is another way to integrate our specialized services with these products.

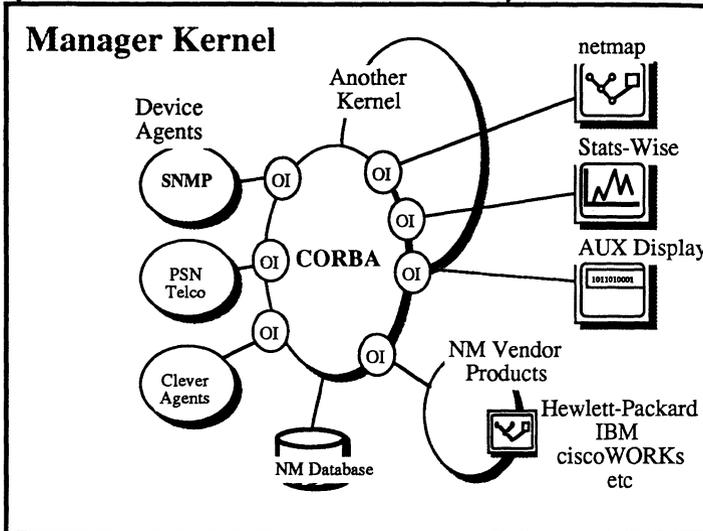
- **Data Management** - Data management lets us tie together the device and presentation agents.

We have chosen the CORBA[5] approach because this has the functions we need and looks as if it will be widely embraced by builders of all sorts of information systems. CORBA is an architectural approach which is supported by many vendors, both large and small. It provides a way to organize data that is needed by multiple users in different physical locations. CORBA provides guidelines for solving a number of difficult distributed operating system problems that we encounter in an NM system, such as providing common interfaces to data repositories and distributing a common consistent view of data to multiple users that are temporally and spatially separate.

5. Future

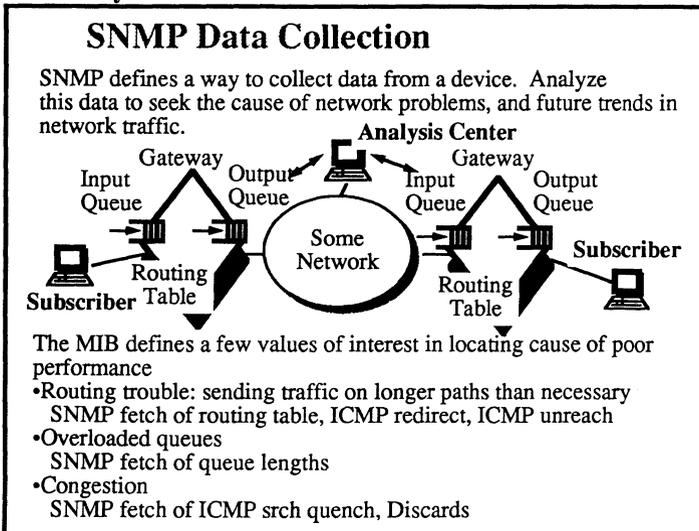
We have our new NM platform working in the lab and we are using it in several research and

development efforts, but much work remains to be done. We are at a very early stage of application of CORBA - since the specification itself is still evolving. However we have a precursor of a CORBA system, called BBN/CRONUS[6] that we are using very effectively to provide a Object Oriented data broker and wide area distribution system. We have fielded an NM



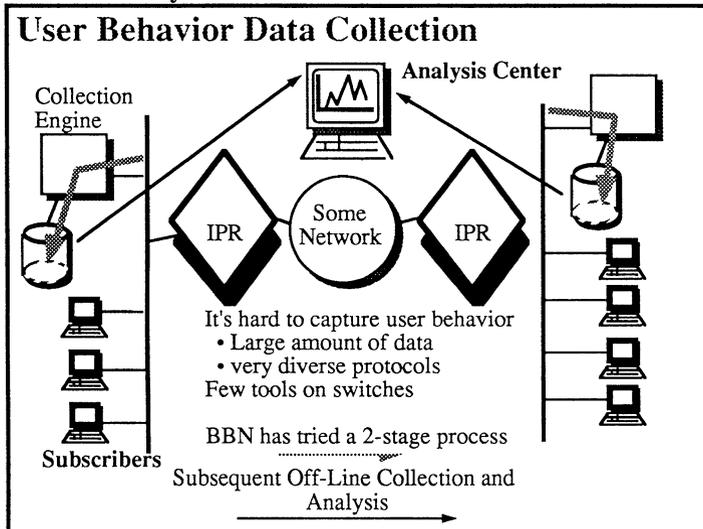
system for testing these ideas and we have interfaced several other commercial and research platforms to it. We are studying the prospects for a commercial release of a product based on this approach.

5.1. Network Analysis



BBN has always had a strong interest in network analysis. Usually, the expense of the networks we build must be justified by cost savings. In most networks, the biggest costs are for the long distance transmission links, and so we seek ways to ensure that all links are being used effectively. If a link is underutilized, its cost is wasted, and if it is overutilized, network performance may suffer. Over the years, we developed many tools for collecting data from an operating network, and analyzing the data to better understand the network's present operation, and to predict future trends in network traffic. SNMP[7] facilitates the raw data collection because this single protocol works with many devices. The associated Structure for Managed Information (SMI) defines a large number of items which are available or items which are available can be collected using SNMP. SNMP with SMI will be adequate to meet statistical collection needs for the foreseeable future.

5.2. Requirements Analysis



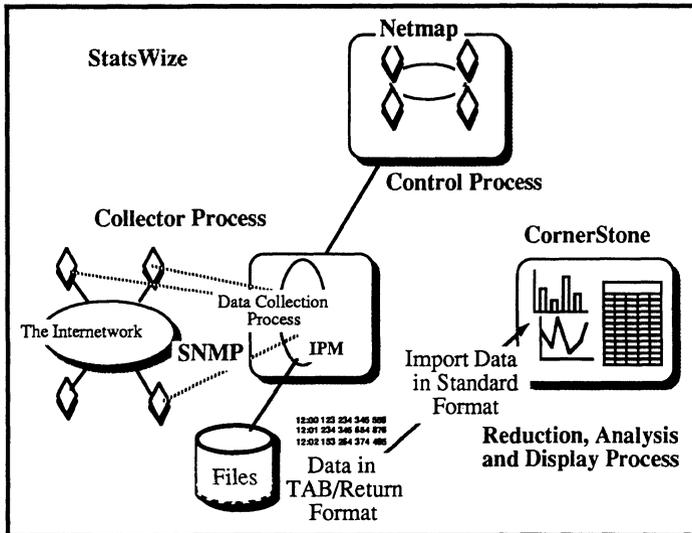
In addition to monitoring the behavior of network devices, we seek to understand the behavior of network users. This is much harder, since users are often independent of network administration, and there is great variety in their computer types and behavior patterns.

One technique which works is to copy all the user packets from each LAN or user interface, into a file and analyze this file off-line. This generates a vast amount of information, most of it not relevant to the analysis, but in the absence of standards for obtaining this data, say from within the hosts, this is the best option. Another technique, called Host Traffic Matrix (HTM) collection also works, but does not collect as much information. An HTM collection consists of records of the number of times a particular source/destination address pair is seen at an interface. These records are generated by some commercial IP routers and packet switches. In either case, the resulting data must be filtered, to extract useful information.

5.3. A Net Analysis Approach

Each network is different and the engineers responsible for operations are often the most suitable candidates to analyze the net's behavior. To support these engineers, BBN has developed a refinement of some of the tools that we have used for many years in the oversight of first the Arpanet and later a variety of X.25 and IP networks. We have combined the data collection, the data management and the data manipulation and display into an integrated tool called StatsWize. Stats-

Wize can be configured to collect any desired set of SNMP variables, from any set of IP devices, according to a schedule determined by the experimenter. The collected data is "cleaned up" and aligned with a common time scale, and then filed in a form which makes retrieval quick and efficient. Various general-purpose statistical analysis software packages can then be used to experiment with the recorded data. A troubleshooter seeking a low-frequency fault might choose one form of analysis, while a network planner might look at traffic flow level trends. By presenting the raw data to the on-site experts, the best of local knowledge and powerful tools are combined.



It is interesting to note that some of the most effective tools for analyzing today's Internets are based on analyses that precede the startup of the first Arpanet node. When these basic mathematical techniques are placed in the hands of the practitioners of network operations and management, they help improve the network availability and efficiency.

Bibliography

- [1] Heart, Frank, et al, "The Interface Message Processor for the ARPA Computer Network," American Federation of Information Processing Societies (AFIPS) Conference Proceedings, Vol 36, p.551-567, June 1970
- [2] Hasegawa, T. "Digital Data Dynamic Transmission Systems," IEEE Transactions on Communications Technology, September 1964
- [3] Baran, Paul, "Reliable Digital Communications Using Unreliable Network Repeaters," National Technical Information Service, NTIS Accession Number AD-616 671.
- [4] McKenzie, Alex, et al, "The Network Control Center for the ARPA Network," Proceedings of the 1972 International Conference on Computer Communications

[6] Schantz, R.E, Thomas, R.H, and Bono, Girome, "The Architecture of the Cronus Distributed Operating System," Proceedings of the 6th International Conference in Distributed Computing Systems, 1986, IEEE Computer Society

Additional Reading

[5] on CORBA, The Common Object Request Broker Architecture:

Betz, Mark, "Interoperable Objects," Dr. Dobbs Journal, Oct 1994 [Better reference]

and, "The Common Object Request Broker: Architecture and Specification," Object Management Group Document Number 93.xx.yy [another reference]

[7] Schoffstall, M. et al, "A Simple Network Management Protocol (SNMP)" RFC1157, at ds.internic.net in /rfc/rfc1157.txt