Intelligent Management for Large Networks

Rajiv Sreedhar, Timothy D. Hill, and Gregory M. Stanley

Gensym Corporation
1776, Yorktown, Suite # 700
Houston, TX 77056

Abstract-- **Enterprise computer networks are rapidly increasing in size and complexity. The typical enterprise network consists of numerous hardware architectures, communication protocols, and user interface technologies. The requirements of high availability and performance place a very high demand on the management of such networks. The ever pressing need to increase efficiency often results in the decline in the network operator to network element ratios resulting in loss of network availability and a decreased level of service. In such scenarios it is essential to integrate an expert system into the network management architecture. The expert system facilitates the transformation of data into value-added information enabling the network operators to efficiently manage the network. In this paper we present details of an intelligent network management system which is successfully being used to manage a large internet service provider network.**

***Index terms*-- Network management, expert systems, correlation, filtering, automation.**

I. Introduction

As corporate computer networks get larger and larger their management gets increasingly more difficult and complex. The number of alarms generated typically grows in proportion to the size of the network. As the network size grows it becomes increasingly essential to incorporate some form of intelligent, real-time alarm filtering and correlation scheme to reduce the amount of alarms presented to the operators. Modern computer networks also typically contain systems which interact and/or depend on one another, for example file servers, mail servers, gateways, routers, and service access clients. It is thus essential to model the interaction between the various components to be able to determine the effect of system components on service availability.

A systematic approach for managing large networks has been presented in [1]. Fault management and correlation issues in network management have been presented in [3, 6]. Expert systems based network and systems management has been presented in [2, 4, 7, 8]. In this paper we present the design and implementation of an integrated availability management solution of large networks using Gensym's G2 real-time expert system and Gensym's Integrity toolkit. Details of modeling the computer network domain in G2 and the methodology for implementing some typical filtering and correlation schemes in Integrity will also be presented. The paper will also present details of the implementation of such a system for managing very large enterprise networks.

II. The Need for Intelligence in Network Management

The ongoing revolution in information technology has resulted in the deployment of very large computer networks. These enterprise networks are constantly increasing in sophistication and size, and are often being deployed on a global scale. The push to incorporate some sort of intelligence in the network management infrastructure of such networks is primarily driven by:

1. The need to continue managing the increasing network size without a corresponding increase in staff.
2. The need to provide a higher quality of service for the users by ensuring that the service level objectives for the network are maintained.
3. The need analyze event storms and determine the root cause of the problem.
4. The need for increased sophistication required for diagnosis, due to more complex systems.

III. Network Management Paradigms

There are three fundamental types of network management paradigms

1. Exception/Event based network management
2. Polling based network management
3. Hybrid exception directed polling based network management

A. *Exception based network management*

Exception based network management relies entirely on status notifications received from the agents monitoring the devices or from the devices themselves. The major advantage of this approach is that is scales nicely with increasing network size and excessive bandwidth is not used for the network management functions. Basing a network management architecture entirely on exceptions is not desirable because catastrophic failure in a network system may not result in an exception if there is no communication path to the management station. Also even when problems are detected, there is often not enough information to pinpoint the exact problem using only passive monitoring of events.

B. *Polling based network management*

Polling based network management relies entirely on polling the network to determine system state. Typically basic status polling in IP networks is performed using **ping**. System health is monitored by the scheduled polling of several parameters. Typical parameters polled include memory utilization, CPU utilization, and error counts. The major drawback of this approach is that a significant percentage of the network bandwidth is used for network management. This network management paradigm does not scale well with an increase in network size. The network bandwidth and the resources required by the polling engine grow rapidly with increasing network size.

C. *Hybrid exception directed polling based network management*

Most network management implementations consist of a mixture of exception based and polling based management paradigms. Typically the status of the network is determined by periodically pinging every device. A few key parameters such as quality of service and performance metrics may also be polled. Additional system parameters are polled only for a short duration when an event/exception is received from a device. The primary advantage of event driven polling is that it facilitates the collection of system data for diagnosis only when there is event/exception generated. This results in a very scalable network management paradigm and has the advantage of being able to specify the parameters polled, and the frequency and duration of polling depending upon the type of event/exception received.

IV. Intelligent Network Management System Architecture

A typical architecture for an intelligent network management system is shown in Figure 1. The system uses a hybrid of the polling and exception based approaches to network management. The key components of the network management system are the polling engine, expert system, archival facility, and the operator interface. Additional intelligence may reside on expert agents not shown or discussed here to simplify the discussion. For simplicity, we also assume only one expert system, although multiple ones are frequently used.

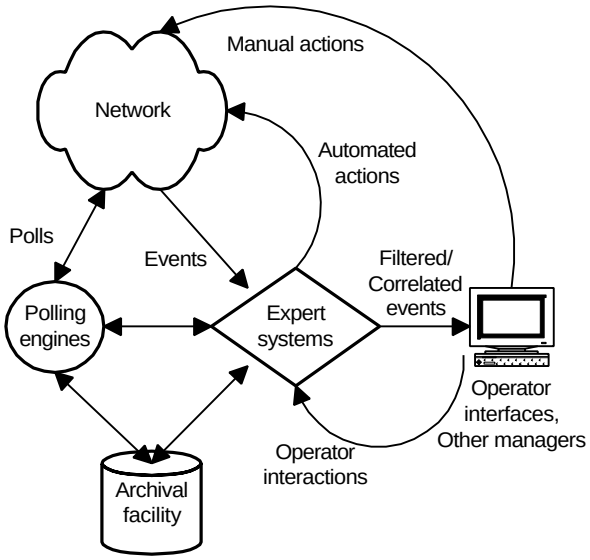

Figure 1: Intelligent network management architecture

The polling engine facilitates the periodic status polling (using *PING* for example) of all the devices on the network. It also facilitates the periodic data collection of key system parameters such as CPU usage, free memory, and various error-counts on a periodic basis if desired. The polling engine generates exceptions for example when a system fails a poll or when certain system parameters violate a specified threshold. The exceptions generated by polling are sent to the expert system for further value-added correlation. The expert system is tightly integrated with the polling engine and has the ability to modify its behavior. The expert system has the ability to modify the list of devices being polled, the polling frequency, the parameters being polled, and the duration of polling.

The expert system filters and correlates all the events generated by the polling engine and those received directly from the network. The expert system is the single point where all the network events are compiled for filtering and correlation. Only filtered and correlated events are passed to the operator interface system for presentation. The ratio of filtered events to raw events depends on the sophistication of the expert system, ratios as high as 1:1000 are not uncommon. Similarly the sophistication of the expert system also determines the quality of the information provided to the network management operator. The expert system is the key component of the intelligent network management system and its characteristics and typical knowledge representation are discussed in detail in the subsequent sections.

The archival facility typically consists of databases to keep a historical record of the system parameters. The historical record is typically used for long term trending and for decision making by the expert system. It is also used for generating historical reports of network utilization and system availability.

The operator interface typically consists of customizable graphic screens to present the network status to the operator. The major components of the operator interface consist of the topological network map representing the connectivity and containment of the network components, browsers to view the events and system parameters, and a number of vendor specific configuration tools.

V. Expert System Characteristics and Requirements

Speed and throughput are the primary requirements for any expert system used in a network management environment. The expert systems have to work efficiently in real-time and process hundreds of events per second. Expert systems also have to process rules, procedures, and user defined behavioral models of the domain, and should also have standards based interfaces to facilitate data exchange with other systems. Some of the typical standards based interfaces are SNMP (Simple Network Management Protocol), TCP/IP sockets, HTTP (Hyper-Text Transmission Protocol), Java RMI (Remote Method Invocation), and CORBA (Common Object Request Broker Architecture).

The expert systems should have an object oriented programming environment to leverage the advantages of object oriented programming. The expert system should also have a user friendly development interface for rapid prototyping and development of the application. Some of the desirable features include graphical objects and a structured programming interface. Traditionally, expert systems have been associated with rule-based systems. Modern expert systems combine many technologies, so that rules are only a small part of the solution.

VI. Knowledge Representation for Intelligent network Management

The typical functions performed by the expert system can be classified into five general categories:

1. Alarm filtering
2. Alarm escalation
3. Alarm correlation
4. Prediction of failure effects
5. Fault diagnosis
6. Mitigation and recovery

A. Alarm filtering

Alarm filtering typically involves drastically reducing the number of alarms which are presented to the network operator without reducing the information content or compromising the ability of the operator to recognize network problems. Alarm filtering forms an integral part of any intelligent network management system because it prevents the network operators from being swamped during event storms.

Duplicate alarm suppression:

The majority of embedded agents on devices are designed/configured to repeatedly send notification to the management system as long as the anomalous condition exists. This is usually the case when management protocols such as SNMP (Simple Network Management Protocol) which do not guarantee delivery of the notification are used for managing the network. In such situations it is not uncommon for the management platform to receive the same notification every minute. In this scenario the expert system presents the operator with the first occurrence of the alarm and suppresses all subsequent occurrences until the alarm is acted on by the operator.

Threshold filtering:

There are many situations where the receipt of a single notification does not signify a problem with the network.

However the occurrence of repeated notifications does indeed signify a problem with the network. The threshold filtering rules in the expert system help distinguish between occasional network glitches and persistent problems. In a typical implementation of threshold filtering the expert system only presents the event to the operator if there have been more than **N** occurrences of the event in a **T** minute time interval. The threshold **N** and the length of the sliding time window **T** is configurable for each alarm/event category.

Topological filtering:

There are numerous failure modes where the occurrence of a failure causes the generation of numerous sympathy events. In such scenarios it is essential the expert system filter out the sympathy events and only present the operator with the root cause of the problem. The sympathy alarms are typically generated by devices and systems which are topologically related (via connectivity, containment, or any other relationship) to the root cause of the problem.

Topological alarm filtering capability is an integral part of any intelligent network management system. This capability is especially needed for managing the state of the art network infrastructures based on the ATM technology. In a typical ATM failure scenario a **card-**

down event would cause several sympathetic **physical-port-down** events from the physical ports contained in the card. Each physical-port-down event may in turn result in **logical-port-down** events from logical ports contained in each physical port. It can easily be seen that if each card had 8 physical ports and each physical port had 16 logical ports, a card down event would result in 136 sympathy events being generated. In such a scenario the expert system would recognize the physical-port-down events and the logical-port-down events as sympathy events and filter them from the operator.

Another typical application of the topological filtering is a scenario where an upstream device failure results in sympathy alarms from all devices downstream of the failure. In enterprise computer networks the failure of a router or a critical serial link may result in several hundred sympathy events (**ping-failures**) form the downstream devices. In such scenarios the expert system recognizes that the sympathy events were generated as a result of the upstream failure and filters them out.

B. Alarm escalation

The expert system also plays a key role in prioritizing the tasks of the network operator. The expert system determines relative severity of each event, and the priority of the events severely affecting the performance or availability of the network is escalated. The escalation rules in the expert system ensure that the operator is promptly notified of critical network problems.

Repetition based escalation:
If certain alarms are received repeatedly the severity of the alarm in increased and presented to the operator. In certain cases an audio alert may also be presented to the operator.

Temporal escalation:
The expert system may reissue an alarm to the operator with increased severity if the original alarm is not acknowledged and acted on in a specified time interval.

C. Alarm correlation

The expert system may correlate two or more alarms to issue a correlated alarm containing additional information. The value added information provided in the correlated alarm is extremely useful for fault diagnosis. In some implementations the expert system may also automatically clear an alarm when a clear event is received.

Auto-clearing:
The expert system may correlate an alarm clear event with a corresponding alarm onset event received earlier and automatically clear the alarm condition. There may be a one to one or one to many relationship between the clearing event and the alarm onset events.

Domain based correlation:
The expert system may also perform correlation based on the characteristics of the various devices in the network and the topological characteristics of the network. The implementation of domain based correlation rules requires an intimate knowledge of the domain. In most cases the correlation rules are non-trivial and difficult to conceptualize and implement.

Correlation and diagnosis based on queries of event histories:
Typical correlation and diagnosis relies heavily on queries of event history. Typically a "list" of target objects is assembled based on relationships to a particular object. Then the query also checks the sending object and the event name (with optional wildcards). The query specifies a pattern of events in related objects.

D. Prediction of failure effects

The expert systems can reason over network behavior and offer expert advice to the network operator. The system can also perform "what - if" scenarios and increase the awareness of the network operator.

E. Fault diagnosis

Expert systems are ideal for implementing complex fault modeling and diagnosis, [12]. Some of the common techniques for diagnosis are rule based, fault tree analysis, causal directed graph analysis, and state transition based fault analysis.

F. Mitigation and recovery

Fault mitigation and system recovery form an important component of any intelligent network management system.
Automation of the routine recovery tasks results in reduced operator fatigue thus improving efficiency. This is especially important for managing a large network since there may be a significant number of such tasks performed.

VII. Intelligent Network Management System Implementation

An intelligent network management system based on the architecture described in Figure 1 was implemented to manage a large Internet Service Provider network. The network consisted of about 125,000 managed entities geographically distributed

across the United States. The network was managed using the Simple Network Management Protocol (SNMP). The use of SNMP for network management is extensively discussed in [9, 10, 11]. The implementation of the intelligent network management system is described in Figure 2.

The management system implementation has three major components consisting of the polling engine, expert system, and user interface. Each of these systems run on separate workstations and communicate with each other using a combination of SNMP traps and remote procedure calls (RPCs).

The expert system component of the intelligent network management system is implemented using Gensym's G2
real-time expert system and Gensym's Integrity layered product for network management. The major components of the expert system consist of a topological domain map capturing the containment and connectivity of the network. The domain representation in G2 is also object oriented and captures the behavioral characteristics of the network elements. All the managed entities in the network are derived form the *MANAGED OBJECT* super class. The managed object class is sub-classed to endow the elements with additional behavior and attributes

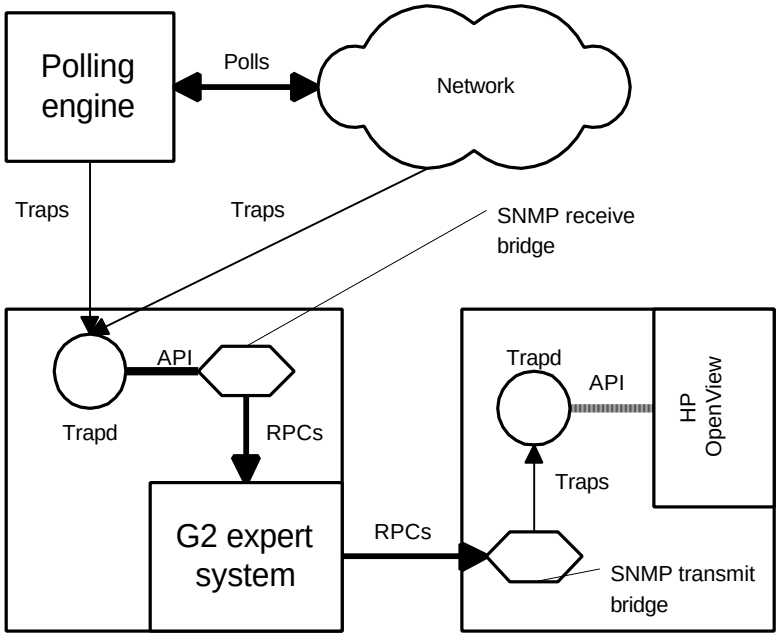

Figure 2: Intelligent network management system implementation

Polling Engine:
The polling engine polls the network using a combination of ICMP pings and SNMP gets. The ICMP ping is used for status polling and determines the availability of the network devices. The SNMP polling is used for determining the performance of the devices on the network. The polling applies simple threshold tests on the data collected and sends a notification to the expert system if any threshold is violated. The notifications from the polling engine to the expert system are facilitated using SNMP traps. The typical parameters collected in polling are free memory, CPU utilization, various error counts, and parameters related to routing. In this architecture the routine task of scheduled polling is off-loaded from the expert system and the expert system is then used only for high value tasks.

Expert System:

specific to the sub-class. A partial domain hierarchy representation of the domain objects is shown in Figure 3.

The filtering, correlation, and diagnosis logic for the network are represented as a combination of rules, procedures, state transition models, and causal directed graph models. The expert system contains event-driven rules which respond automatically on receipt of new data and are used to invoke other rules, methods, or procedures.
The expert system also has a (short-term memory) facility of archiving events for temporal correlation.

The expert system receives the notification of network events using SNMP traps. The SNMP (receive) bridge process shown in Figure 2 receives the traps from the trap daemon via an application programming

interface, the bridge then passes on the information contained in the traps to the G2 expert system using remote procedure calls. The expert system processes these events and passes the correlated events to the operator interface using SNMP traps.

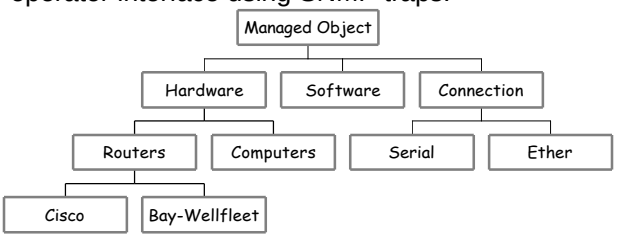

The network management operator interface is a customized Hewlett-Packard OpenView Network Node Manager (OV NNM) application. The OpenView NNM provides the operator with graphical view of the network. The NNM is a point of integration for the various enterprise specific element managers such as Cisco Works.

The NNM application also has a customized event browser for viewing the network events and also a MIB browser for viewing the management information base of the various agents.

VIII. Operational Details of the Intelligent Network Management System

The intelligent network management system has been implemented for managing a large internet service access network by one of the worlds leading telecommunication service providers. The network consists of over 125,000 managed entities. The event rates in the network are summarized in Table 1.

Table 1: Network statistics

Avg. # of raw events/day	800,000
Avg. # of correlated events/day	1000
Sustained raw event rate	8 to 12 events per second
Peak raw event rate	50 - 100 events per second

IX. Conclusion

The intelligent network management system presented has been successful in managing a large industrial strength network. The network management system has been in operation for about four years. The availability of the network increased substantially as the intelligence embodied in the expert system evolved in sophistication. In the absence of the expert system the network management team cannot handle a load of 800,000 traps a day. This is a success story of how an expert system transforms data into information enabling the network operation group to efficiently manage a large network.

References

[1] Ananthanpillai, R., *Managing Messaging Networks a Systematic Approach,* Artech House press, Boston, 1995.

[2] Ball, L., *Network Management with Smart Systems,* McGraw-Hill, Inc., 1994.

[3] Byrne, C., "Fault Management", in *Telecommunications Network Management into the 21st Century,* ed. S. Aidarous and T. Plevyak, IEEE Press, 1994.

[4] Cronk, R., Callahan, P., and Bernstein, L., "Rule-Based Expert System for Network Management and Operations: An Introduction", *IEEE Network Magazine,* September 1988.

[5] *G2 Reference Manual*, Ver. 5.0, Gensym Corporation, 125 CambridgePark Drive, Cambridge, MA 02140, USA.

[6] Jakobson, G. and Weissmann, M, "Alarm Correlation", *IEEE Network Magazine,* Vol.7, No. 3, November 1993.

[7] Muralidhar, K., "Knowledge-based Network Management", in *Telecommunications Network Management into the 21st Century,* ed. S. Aidarous and T. Plevyak, IEEE Press, 1994.

[8] Rabie, S., Rau-Chaplin, A., and Shibahara, T., "DAD: A Real-Time Expert System for Monitoring Data Packet Networks", *IEEE Network Magazine,* September 1988.

[9] Rose, M. and McCloghrie, K., *How to manage your Network using SNMP: the network management practicum,* Prentice Hall 1995.

[10] Stallings, W., *Network Management,* IEEE Press, 1993.

[11] Stallings, W., *SNMP, SNMPv2, and CMIP: the practical guide to network management standards,* Addison-Wesley, 1993.

[12] Stanley, G. and Vaidhyanathan, R., "A generic fault propagation modeling approach to on-line diagnosis and event correlation", *Proceedings of the IFAC workshop on On-line Fault Detection and Supervision in Chemical Process Industry,* Lyon, France, June 1998.