I. Introduction

The term "proactivity" has come into much wider use in the last several years. For something to be "proactive" means that it is capable of dealing with problems by anticipating and preventing them, not merely by reacting to problems after they have occurred. In today's increasingly complex business and technical environment, with needs and demands for better support growing constantly, reactive support has been taken almost to its limits. Proactive support is the next step towards providing the highest possible level of service.

Support has historically been divided into the categories of hardware and software, and has been packaged and sold as such. In the last few years, as the use of networks has grown, new ways of looking at support have emerged. Network support has become an entity in its own right; it is not merely the sum of existing hardware and software support, but is the product of a new, more integrated approach.

HP's network support strategy has many facets, but this discussion will focus on only one: HP's strategy for network support tools. These tools are used by internal HP personnel in support of customer networks. The strategy used in design and development of these tools has three major components.

The first component of the strategy is remote support. HP's goal is to provide as much support on a remote basis as possible, in order to take advantage of centralized expertise of support personnel and to eliminate travel time to a customer site. The second component of the strategy is a focus on configuration. Configuration of all network elements, whether hardware or software, is a frequent source of problems with a network. Focusing on simplifying, checking and correcting configuration is an efficient way to solve a large number of network problems. Finally, the third component of HP's network support tools strategy is an emphasis on proactivity. As with hardware and software, proactivity and proactive tools for network support are the next step in providing customers with the higher levels of overall support which are required today.

This paper will discuss how HP is putting these objectives into action by developing a tool for proactive network diagnosis. It will describe in more detail why proactive network diagnosis is needed, how it operates, and what information it provides.

II. Purpose

The need for proactive network diagnosis grows out of the nature of networks. Networks are intrinsically more complex than either software or hardware, since they contain elements of both. Their many components require very specific configuration, and the configurations for the different components are highly interdependent. In addition, the networks themselves must be configured in terms of the routes, addresses, and capacities which characterize them.

The trend in networks, as in computer systems themselves, has been a steady increase in size, speed, and complexity. This trend is even more marked in networks, as the ways in which people use computing have changed. Users are not only performing more of the same activities faster, but also performing different activities. The shift towards distributed computing operations and towards the use of interconnected personal computers and workstations represents a change in what is done as well as how fast it is done. This increasing use of distributed computing creates a corresponding increase in the need for

network support. Network operations in general have become a more critical part of users' businesses, and this trend will continue. Networks are also incorporating more multivendor devices, as open standards become increasingly accepted. A better way to support these ever more critical, ever more complex networks is needed.

Proactive network diagnosis is such a way. In general, "proactive network diagnosis" means monitoring events on a customer's network and identifying problems before they occur. Specifically, proactivity is provided by programmatically reading the network log files generated by the various levels of networking software and hardware, and analyzing the events and event rates recorded in the logs. The customer's system, with its network log files, serves as a "window" into the network. When the events and event rates found in the log files are analyzed, thresholds for acceptable event rates are used to determine when an exception requiring attention has occurred. Thresholds are set at a level which allows exceptions to be detected before a network goes down or becomes noticeably impaired. When an exception which can be resolved through customer action is detected, it is reported to the customer. If an exception requiring HP support personnel action is found, information about it is electronically transmitted to HP's Response Center, where the appropriate personnel can take action.

The need for this type of support is apparent. The goal of all support is to maximize the product's - in this case, the network's - usability and usefulness. This includes uptime, throughput, and reliability, among other factors. To provide this, problems need to be detected and resolved before they negatively affect customers' normal operations.

III. Operation

In order to understand how proactive network diagnosis operates and what it can provide, an understanding of the framework within which it fits is required. HP's offering in the area of proactive network diagnosis, known informally as Network Predictive, builds on the existing structure of the Predictive Support product. Predictive Support can be thought of as proactive system diagnosis. It has been running on customer HP3000 systems for several years, providing proactive support for disc drives, tape drives, and memory. Network Predictive adds network software and hardware support to the HP3000.

Products supported through Predictive Support or Network Predictive must have a high degree of internal error detection and reporting. For each product selected, product experts from the Response Centers and from manufacturing divisions model the degradation and failure modes of the product. The results are reduced to a set of rules which are incorporated into Predictive Support or Network Predictive. Each rule describes an event and specifies the frequency at which its occurrence is considered to indicate the existence of a problem; this frequency of occurrence is called a threshold. For example, a rule might be written for the network transport subsystem which describes the event "store and forward packet discarded" with the threshold "10 in a day". This would mean that, if ten or more store and forward packets are discarded in a day, Network Predictive would notify the Response Center that a potential problem has occurred. The definition of the rules and thresholds is a dynamic process involving feedback and adjustment, so the Response Center experts must monitor the effectiveness of the rules established for each product.

After Network Predictive is distributed to the customer system, analysis begins. The Predictive job is normally run nightly, during off-peak processing hours. Network Predictive executes in four basic phases. First, log files are scanned and error data is collected. Second, the error data is expressed in a generic format. Third, trend detection is

performed, and finally, if necessary, the appropriate actions are taken to solve any problems. This may include automatic generation of a message to the console operator indicating what actions the customer may take to resolve the problem, or the generation of a call to the Response Center. Overall processing is controlled by the Predictive monitor process. Figure 1 shows a graphic representation of the Predictive (including Network Predictive) architecture.

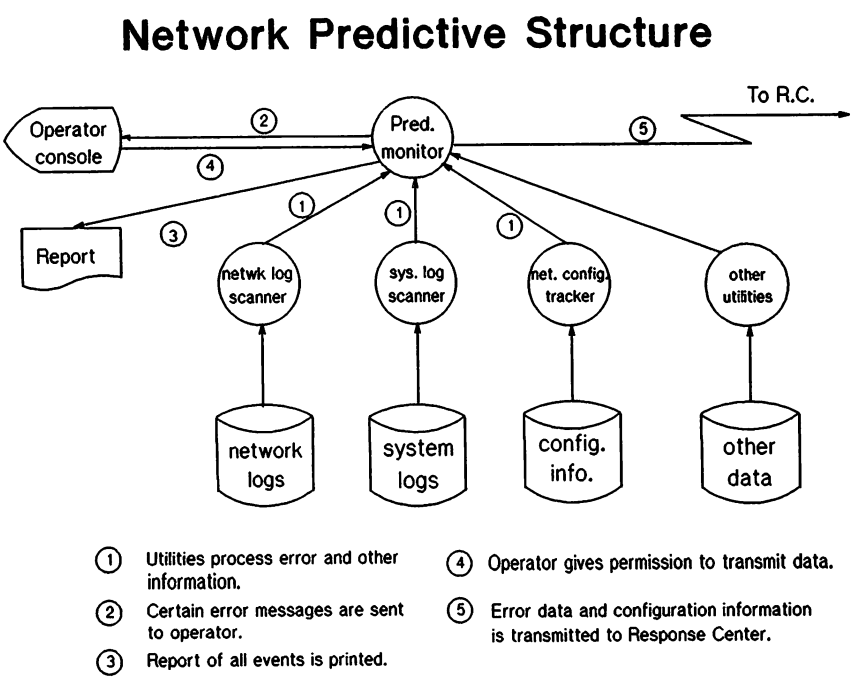

Figure 1: Structure of Network Predictive

Network Predictive uses special utility programs to collect error data. Each utility is launched as a child process of the monitor and retrieves the error data for a specific class of products, using built-in information about the format in which that class of products performs error logging. In the second phase of processing, the utilities translate the many different error data formats into a common message format, so that the Predictive monitor can use the same trend detection algorithms to determine whether a failure is imminent on the specific system component, regardless of which type of component generated the error. The messages contain information identifying the product and the type of error involved.

When the messages are received by the Predictive monitor process for third- phase processing, the error data is passed through a trend-detection algorithm. If the results indicate that an undesirable trend has been established, the appropriate action is triggered.

For trend detection, the frequency of occurrences of significant events must be established. This is done by tracking the number of occurrences of an event and weighting it by a factor such as time over which the number of events occurred, or number of related normal events (e.g. number of packets sent). In addition, since simple continuous tracking of abnormal event occurrences and weighting factor would lead to a dilution of the statistics, the accumulation of the weighting factor must be limited to create a sample across which the frequency of occurrence can be analyzed. When the trend detection algorithm detects that the number of occurrences of a given event, divided by its weighting factor and taken across the designated sample size, exceeds the threshold defined for that event, an action is triggered.

The fourth and final phase of Network Predictive processing is taking the specified action. Depending on the rule for the given event, the Predictive monitor will either send a message to the console to inform the operator about the problem, or it will use a communication link to transfer the information directly to the Response Center for investigation by support personnel. This information is then supplied to the customer in the form of a printed report.

The data communication link between the customer HP3000 system and the Response Center uses a remote support modem installed on the customer system. At the end of the nightly Predictive run, Network Predictive uses the modem link to transfer any event information gathered during that run to the Response Center. This transmission may be accomplished completely automatically, or it may be configured to require operator intervention for security reasons.

When the data reaches the Response Center, an event generation process loads it into a Response Center database for later examination by Response Center personnel. The process also generates actual Customer Service Orders within the Response Center's call tracking system, insuring that each call will receive prompt and thorough attention, including remote handling and field referral if necessary.

Network Predictive provides the capability for proactive network diagnosis. The question then becomes one of content: what information is available for predictive analysis?

IV. Network Information

Network Predictive adds a log file scanning utility to those already incorporated in Predictive Support. This new utility operates on the log files generated by HP's Network Services (NS) software. New rulesets for analyzing the data contained in the network log file have been developed. Rulesets have also been added to the set of rulesets currently used in analysis of MPE system log files; the new rulesets contain rules for analyzing the operation of INP (Intelligent Network Processor) devices. Other differences between Network Predictive and the current Predictive Support products include the increased use of rules which use a number of normal event occurrences (units) instead of time as a weighting factor, and Network Predictive's far greater number of events which generate messages to the console operator rather than to the Response Center.

Although HP's NS software was not designed to be in accord with the OSI (Open Systems Interconnect) model for standardization of networking software, the entities which comprise it can be grouped to approximate the 7-layer OSI model. Figure 2 shows the NS entities and how they fit into the OSI model. The rulesets developed for the log files generated by the NS software were chosen to operate on entities at the network, transport, and session

layers. The software at these layers logs information with the most appropriate amount of detail and usefulness. Entities currently covered by rulesets include TCP SIP, TCP PM, IP, and NetIPC.

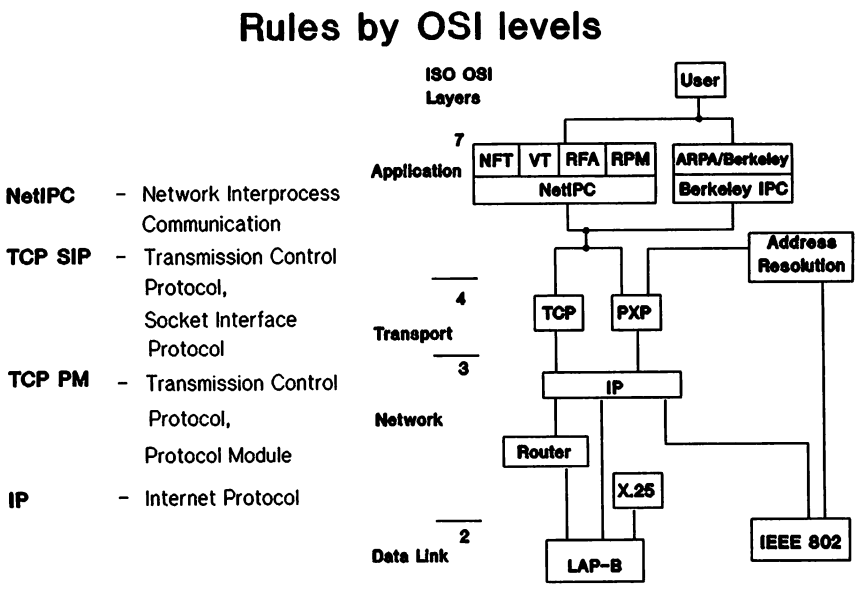

Figure 2: NS Entities in the OSI Model

The rules for the NS software logs and for the INP device information logged in the MPE system logs were developed using the technique known as "knowledge engineering". Experts in the various subject areas for which rules were being developed were identified. Development engineers working on the Network Predictive product then spent time with these experts, extracting knowledge from them and translating that knowledge into rulesets suitable for use by the Network Predictive event data collection and analysis processes. Most of the initial development of the rules was done by working with lab personnel, especially those who had been involved with the initial development of the product/entity. The rules were reviewed and the thresholds established or refined by field personnel, particularly those from the Response Center as well as the Network Marketing Center and the Customer Escalation Centers. This approach combined a thorough knowledge of product internals with practical experience of how often events are likely to occur and how serious they are. The field personnel also helped to develop the troubleshooting guidelines defined for the different events.

In their final form, the rules for interpretation of network events can be divided into five main categories: line and transmission errors, configuration errors, corrupt or defective

software, corrupt data files, and hardware errors. The category of line and transmission errors includes such events as packet retransmissions and checksum errors, and usually indicates modem or line problems. Configuration errors fall into several subcategories, including resource allocation problems, incorrect address configuration, and inadequate system table allocation. These errors can normally be resolved by the customer; documentation accompanying the console messages generated by Network Predictive describes what steps the customer can take to solve the problem. Errors in the category of corrupt/defective software indicate that a software module has either become corrupted and must be restored, or contains a software defect which must be resolved by factory personnel. INP problems can also be caused by corrupt INP download code. Corrupt data files are usually network configuration files and must be restored from good versions. Finally, hardware errors indicate INP problems.

Not all of these categories are truly "proactive" in nature. The line and transmission errors, configuration errors, and INP hardware errors are the most likely to indicate impending problems. The other categories are included for thoroughness. While many of those errors are likely to show up to the customer before they are ever seen through reading a log file, some of them may pass by unnoticed, or with only a temporary workaround action being taken. Inclusion of these rules insures more complete coverage for the network.

Another important part of Network Predictive which will not be discussed in depth here is its implementation of network configuration tracking. Another utility process running under the Predictive monitor collects network configuration information from its various sources on the customer system, compares the information to the configuration information from the previous night's run, and sends the most up-to-date version through the datacomm link to the Response Center. Only changes are sent across the link once the initial complete configuration has been sent. When the information arrives at the Response Center, it is loaded into a database for use by Response Center personnel, allowing them quick access to configuration information while they work to resolve a proactive or reactive call. The Predictive Support product will be adding similar utilities in its next release, to track system configuration and software version information.

The usefulness of the network information provided by Network Predictive is dependent on the usefulness of the information logged by the network software and hardware elements. Designing for supportability is a key to successful proactive network diagnosis in the future.

V. Conclusion

Proactive network diagnosis is the next step in HP's network support tools strategy, which itself is part of HP's network support strategy: to maximize the use that customers can get out of their networks through increased reliability, increased performance, and reduced cost. By gathering the error information contained in network logs, processing it through the collected knowledge of experts, and taking immediate action on any impending problems that are detected, the Network Predictive implementation of the proactive network diagnosis concept helps to fulfill the goals of HP's network support strategy and to meet today's changing network support needs.