Remote Configuration Tracking: The CONFIRM Concept

Robert Poling

Product Support Division
Supportability Methods Lab
100 Mayfield Avenue
Mountain View, CA 94043

Introduction

CONFIRM (CONFiguration conFIRMation) is not a product. Nor is it the title of a project. CONFIRM is a support philosophy: to provide HP's support organization with *centralized* access to configuration information for remote customer systems.

Problem resolution has traditionally been viewed as a support service which applies to a system or group of systems belonging to a single vendor. When problem resolution required configuration information from the customer system, the support engineers had the following options. Depending on the level of sophistication required to get the data and the amount of data needed, they could request data from the customer. Large volumes of data would have to be transmitted or mailed to the support engineer. Alternatively, an on-site visit by a field engineer could be scheduled -- costly in both time and money. The increasing size and complexity of systems and the growing use of networks placed increasing importance on the availability of configuration information and required new support strategies for problem resolution.

Major developments in HP's support strategy include a focus on proactivity, remote support, and configurations. First of all, in support, *proactive* means dealing with problems by anticipating and preventing them, rather than reacting to problems after they have occurred. Second, 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 minimize travel time to a customer site. The emphasis on proactive diagnosis means monitoring events on a customer system and identifying problems before they occur. Third, there is a major focus on configuration since mis-configuration of hardware or software is a frequent source of problems, especially in networking environments. Configuration simplification and validation is an efficient method of avoiding or solving a large percentage of system and network problems.

CONFIRM is a natural extension of these support strategy developments which HP is pursuing by implementing tools for *remote configuration tracking*.

The goal here is to examine CONFIRM both as a conceptual development in support and as it is being implemented at HP. This paper will discuss in detail why remote configuration tracking is needed and what its benefits are. It will also describe the operation of existing remote configuration tracking tools and how CONFIRM fits into HP's support process.

Purpose

Remote configuration tracking has four major functions: retrieval, change detection, reporting, and centralization. The *retrieval* function performs the task of collecting configuration information from various sources on the customer system. The techniques used to extract and format configuration data depend on the source and may include low-level access to system data structures, database queries, or execution of another process to generate configuration data. The *change detection* function monitors the configuration for any changes which have occurred. Each change is noted as a configuration *event*. The *reporting* function handles the transfer of configuration events from the customer system to a central site via a data communications link, where the events are used to maintain up-to-date configuration information. And most importantly, the *centralization* function facilitates access to the data by support personnel. The user interface provides support engineers with a coherent, standardized way to access widely varied types of configuration data and multiple configurations.

Remote configuration tracking has evolved from the need for improved troubleshooting, timely access to configuration data, and access to multiple configurations.

Historically, configuration information has been collected from customer systems by remote support personnel on a *reactive* basis; that is, as needed. Interactive remote support has improved the situation by providing faster access to data and reducing the demands on customers in the support process. However, the effectiveness of interactive remote access is limited by slow data transfer rates, restricted system accessibility, and cumbersome user interfaces for data retrieval.

These difficulties are further exacerbated by the rapidly increasing size, speed, and complexity of computer systems and computer networks. Larger and more varied applications, distributed computing, multivendor networks and increasing user sophistication all contribute to the growing importance of configurations in support. As a result, problem resolution requires configuration information which must be retrieved from multiple sources on one system, must be current since configurations are increasingly dynamic, and must be available from multiple systems when data communications is involved.

Remote configuration tracking provides centralization for customer system configurations using an automated process. It gives support engineers fast access to formatted configuration data which is accessible regardless of customer system availability.

When configurations are tracked remotely, the customer benefits directly from increased system uptime. It takes HP less time to solve problems when they do occur, whether detected reactively or proactively. And, the customer is less burdened by involvement in the support process.

Operation

The discussion on operation of configuration tracking is in two sections. The first section is an overview of the design and data-flow of remote configuration tracking tools as they have been implemented at HP. The latter section is a survey of existing implementations which will provide more detail on the types of configuration data tracked by each tool.

Overview

HP's Predictive Support provides the support organization with a proactive tool for detecting potential problems before a system's operation is noticeably affected. It is provided as part of HP's Standard, Basic and Guaranteed Uptime support contracts. In addition to its proactive function, the A.02 release of Predictive Support includes several utilities which provide remote configuration tracking for the HP3000 family of computers. These utilities share a common functional architecture which can be divided into three components: onsite, datacomm, and Response Center. The *onsite* component performs the functions of retrieval and change detection. This component represents one of the Predictive tracking utilities. The *datacomm* and *Response Center* components perform the functions of reporting and centralization, respectively. These components are provided by existing Predictive Support software. An operational overview for a Predictive tracking utility follows a sequence of steps which is illustrated in Figure 1.

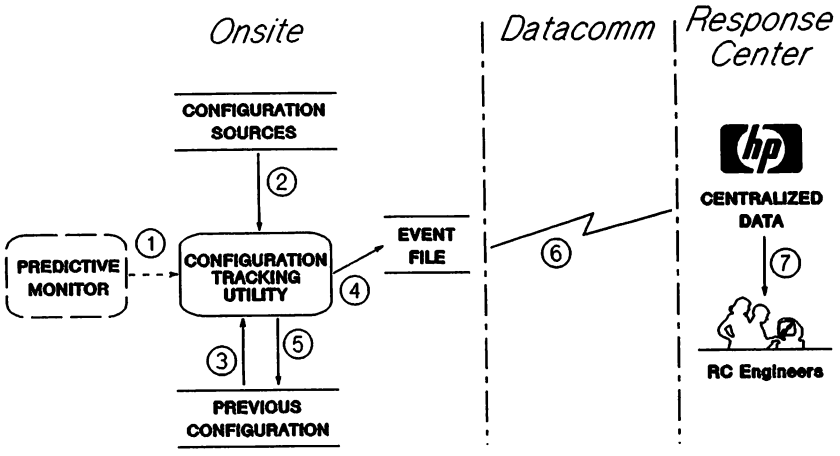

FIGURE 1 – Operational Overview of Predictive Tracking Utilities

1. Predictive A.02 is run every twenty-four hours on a customer's HP3000. Each time Predictive Support software runs on the customer system, it executes the tracking utilities as child processes.

2. The tracking utility collects data from the system sources which contain the current configuration information. If necessary, the raw data is parsed, type converted, and/or reformatted.

3. The current configuration is compared to the contents of the previous configuration file which was kept by the utility during its last run. Note that non-configuration information (current date, number of accessors, etc.) are masked or ignored to eliminate unnecessary configuration events.

4. Any differences in configuration are timestamped and reported to the event file. If the previous configuration file does not exist (as for a first run), it is created and the entire configuration is sent as events to the event file.

5. The previous configuration file is updated to reflect any changes in the configuration; new or modified entries are timestamped.

6. Software on both the customer system and the Response Center system is required for the datacomm component. The datacomm software on the customer system establishes a modem link to the Response Center System. Events generated by the onsite software are transferred across the datacomm link to the Response Center system. The datacomm software on the Response Center system verifies the integrity of the transferred data and stores it.

7. The stored data are then processed by Response Center software which uses the

events to update the configuration in a knowledge base. This component also provides the interface through which Response Center Engineers access the knowledge base.

These Predictive utilities all follow the same basic steps for centralization of configuration information. The major differences between them are the type of configuration information they monitor and the access methods they use to collect that information.

Existing Implementations

Three configuration tracking utilities are being released with Predictive A.02. These utilities monitor a customer's network configuration, system software configuration, and hardware configuration, and are called NETTRACK, SWTRACK, and SYSTRACK, respectively.

The network configuration utility, NETTRACK, gathers its information from a number of different sources and subsystems on the HP3000. Figure 2 illustrates how NETTRACK processes network configuration information and generates network configuration events.

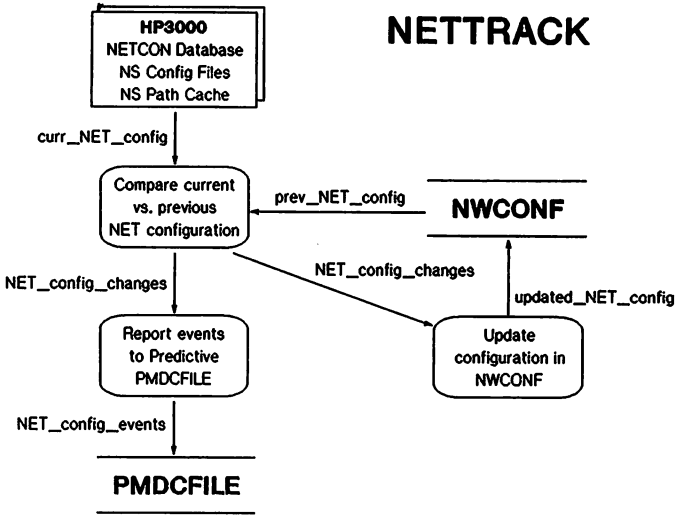

FIGURE 2 - NETTRACK Data Flow Diagram

When NETTRACK is activated it collects the current network configuration into a temporary file. It then compares the current configuration to the previous configuration in data file NWCONF for any differences. Any changes which are detected will be reported as configuration events to the Predictive event file, PMDCFILE. The temporary file then becomes the new NWCONF file. If NWCONF did not exist (as for a first run), the temporary file becomes NWCONF and all data are processed as configuration events. Predictive datacomm will handle the transfer of events to the Response Center.

This information is collected every time NETTRACK is run and will range in size from approximately 80 to 350 records. The number of events generated will depend on how dynamic the network configuration is.

NETTRACK collects and maintains the following information:

- DS X.25 Remote Node information, which includes node name, PDN address and phone number, and additional information obtained from the NETCON database.

- DS X.25 Line Characteristics, which includes logical device number, local PDN identifier, name and address, level 2 and level 3 parameters. This information is also obtained from the NETCON database.

- Network Directory information, which includes global and directory entry information from the Network Directory file NSDIR.

- LAN and Router remote node information, which are derived from NS path cache reports and maintained over time to provide an approximate remote node topology for these networks.

- NS/3000 Critical Summary information from resident NS configuration files (NMCONFIG, NSCONF, and DADCONF).

The second configuration utility, SWTRACK, keeps track of information for installed software program files, SL files, and SL procedures. The operation of SWTRACK is similar to that of the other utilities, with the addition of *maintenance* files. The CONFBILD maintenance file contains the list of software modules to be tracked. For program files and SL procedures, CONFBILD also specifies how and where to obtain associated software versions. For each SL file being tracked, SWTRACK generates a maintenance file which contains SL segment information.

The following information is monitored by SWTRACK:

- Program file information including filename, product number, version (vuf), actual and computed checksums.

- SL file names and segment information including segment name and length, checksum when segment was prepped into the SL, current checksum, and date of last change to entry.

- Name and version of SL procedures.

The third utility, SYSTRACK, covers hardware configuration. The main sources for this information are system memory, system configuration files (CONFDATA and DEVDATA), and the Predictive configuration file.

SYSTRACK maintains the following hardware data:

- System specific information including MPE system id, MPE release id, CPU type, memory size, basic table sizes, scheduling and spool information.

- Volume information for each system and private volume including logical device number, volume identification information, and virtual memory information.

- Device class information for each device class including class name, access type, and logical device numbers for each class.

- Logical device information for each ldev including ldev and DRT numbers, mode flags, class information, and other device configuration values.

- CS information for the communications system and CS details for each CS device including ldev number, port mask, protocol, modes, timeouts, speed, buffer sizes, and driver options.

Since these utilities run within the framework of Predictive Support, customer security is protected by existing Predictive security.

Strategy

Consider then, how remote configuration tracking fits into HP's overall support process.

When a potential problem is detected proactively by Predictive software, it is automatically reported to the Response Center; a problem being handled reactively will be manually called into the Response Center by the customer. Once the problem is classified, a Response Center Engineer (RCE) will determine what, if any, configuration information is needed to aid in problem analysis.

For example, a problem is detected in the X.25 communications between two systems. First the network configuration is checked for compatible addressing, window sizes, etc. Depending on this comparison and other reported symptoms, the RCE may then look up the software configurations of the two systems to check for incompatible datacomm software versions, or the RCE may examine the hardware configurations to verify that the INPs (Intelligent Network Processors) being used for the data link are configured correctly.

In the past, the RCE collecting all of this data would have been required to use 1200 baud remote data links to interactively access the customer systems. Assuming that the RCE could access the customer systems, it would take 30 minutes to an hour per system. But now, the RCE can access multiple configurations through a single interface, at direct connect speeds.

In another scenario, a customer reports intermittent data transfer errors on a LAN which have been occurring for the last week. Before resorting to exhaustive datacomm debugging techniques, the RCE could scan the configurations of systems on the LAN to check for any notable configuration updates which are timestamped as of one week ago. This data could be used to focus problem resolution.

These examples are somewhat simplistic, but they illustrate the usage of configuration information during problem resolution and the usefulness of remote configuration tracking in the support process.

The CONFIRM concept places configuration data at the support engineer's fingertips when

it is needed, in a format that is easy to use. It is not a new support strategy, but a refinement and enhancement of HP's existing approach which integrates hardware and software support with a full range of support services, from consulting services to problem isolation on multivendor networks.

Future implementations of CONFIRM can be targeted for additional systems: dedicated hardware devices as well as general purpose computers. The localization of configuration information will also facilitate development of tools to automate configuration validation and comparison. Using artificial intelligence techniques, tools can be developed to proactively detect potential configuration problems and possible solutions.

Conclusion

HP is dedicated to helping customers maximize their system availability and performance while minimizing their operating and maintenance costs. The remote configuration tracking tools presented in this paper achieve these goals by centralizing customer configuration information where knowledgeable support personnel can access and analyze it, providing timely response to customer support needs.

The CONFIRM concept is a foundation for future directions in support technology and an integral part of HP's support strategy.