



hp data sheet



hp open database software

Open Database software, a component of the HP Performance Management Bundle for NonStop servers, is a powerful performance database.

features at a glance

- Performance analysis with SQL and other powerful query tools
- True network performance database
- Self-describing and extensible
- Open platform for performance applications
- Performance Management Console for collecting and managing metrics



HP Open Database (ODB) is an Open Database Connectivity (ODBC)-compliant SQL database implemented as a Microsoft® Access 2000 database residing on a Windows system-based workstation. ODB software contains detailed performance data collected on HP NonStop servers from HP Measure software, the file system, and the NonStop operating system by the Tandem Performance Data Collector (TPDC), another component of the Performance Management Bundle. ODB software supports both the latest HP ServerNet architecture and the Dynabus architecture.

Much more than a Measure database on a workstation, ODB software contains virtually all detailed Measure data, but the data has been extensively reorganized, summarized, and cross-linked so it can be used effectively at many levels by performance experts and nonexperts alike. For instance, the ODB Process Table contains not only Measure process data but also extended summaries and information from non Measure sources, including process disk I/O, process-to-process traffic, device I/O activity, and system process and process pair information. To obtain the same information from Measure software would require detailed knowledge of multiple Measure entities and extensive computation. Some information (such as process pairs) is not available from Measure software.

performance analysis with SQL and other powerful query tools

ODB software opens the analytical power of SQL to performance analysts. It provides the powerful Microsoft Access query-by-example (QBE) facility that generates SQL queries from a simple drag-and-drop grid. For instance, to see process data along with details of the files a process opens, just place the ODB Process and DiskFileOpen Tables on the QBE grid, drag the Process field in the Process Table to the Process field in the DiskFileOpen Table (the opener process), select the desired fields to display, and click the Run Query button. With practice, far more complex analysis becomes nearly as simple as this example.

true network performance database

ODB software can contain simultaneous samples from multiple NonStop server nodes and can contain a sequence of multiple-node samples over time. These samples are true network samples. Remote opens are cross-linked, so for the first time both ends of intersystem activities can be analyzed together.



self-describing and extensible

ODB software contains detailed descriptors for each of its public tables and fields, including table relationships, field units and rating factors, and detailed and high-level help text. Through database tools, end users and performance application designers can define their own self-describing database extensions.

open platform for performance applications

ODB software is designed throughout for building performance applications. The extensively cross-linked organization and true network character of ODB software simplify applications and make previously impossible ones feasible. Windows and Access technologies together make a large variety of powerful application

builders available, ranging from forms and reports within Access to remote clients that support the ODBC standard. The self-describing nature of ODB software allows applications to be partially or largely independent of the database structure, and its extensible nature allows applications to integrate data from other sources.

performance management console

ODB software now includes the Performance Management Console (PMC), which provides an easy-to-use graphical interface. Designed in a client/server architecture, PMC facilitates the task of collecting and managing NonStop server performance metrics. The client side of this application is based on the Windows operating system and

communicates to a single module running on the NonStop server. A separate host module (PMCMON) carries out user requests and interacts with host-based collector subsystems.

Prior to PMC, a user had to interact directly with different and separate components to achieve the same result. Starting a collection had to be initiated at a TACL prompt, and users had to monitor collections and keep records of collected data manually. The second step of starting deferred processing also had to be initiated manually, followed by yet another manual use of File Transfer Protocol (FTP) to transfer the collected metric file to the workstation. Once the file was on the workstation, the user had to keep track of all current and previously collected and loaded data files before using an Access based interface to load the collected data. Starting client applications was often a separate process, and each client would require the user to locate and open individual samples before analysis could start.

PMC offers a centralized control and management interface to accomplish all of these tasks easily. Using a multitab sequenced approach, PMC guides the user through each required step.

scheduling collections

PMC allows you to schedule collections of metrics on one or more NonStop servers. You set the dates and times, as well as the system or the Pathway server for which the performance metrics should be collected. You also choose where the data will be stored and when it should be processed (see figure 1).

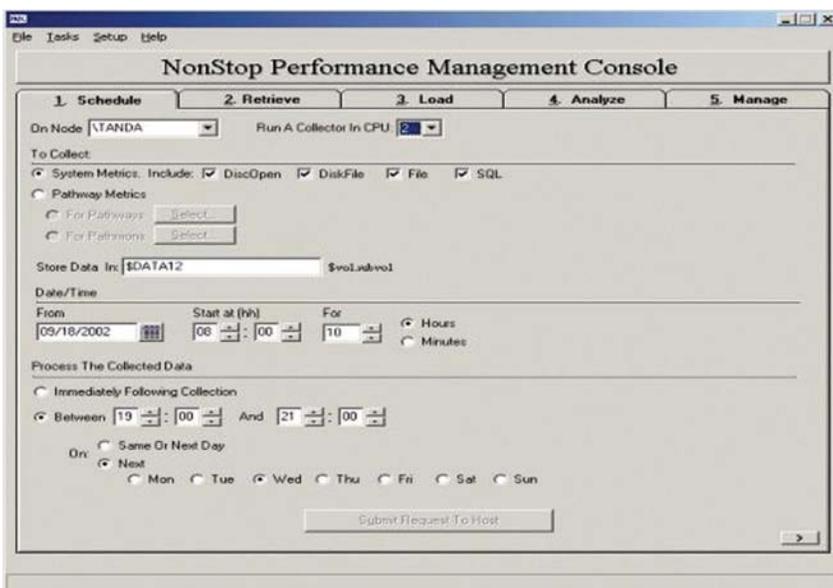


Figure 1. The Schedule window of the PMC interface

retrieving collections

The Retrieve window of PMC lists all scheduled collections on a NonStop server. Several selection criteria options enable you to set parameters where only a subset of all recorded or active collections are displayed. For collections that have completed post processing, you can download the recorded metrics to a data file on your workstation (see figure 2).

loading data files

You can select a UWS data file from the files already downloaded to the workstation and load it into an ODB database. PMC shows a list of both raw UWS data files as well as ODB files on your PC. To further assist you, PMC offers a set of options to locate the proper UWS and ODB. Once the data is loaded, you can use one or more client applications, such as Data Browser, Insight, or Pathway View (also components of the Performance Management Bundle), to analyze the collected data. PMC automatically locates a list of performance client applications and displays them along with a link on its window (see figure 3).

analyzing data samples

From the Analyze window of PMC, you can select one or more data samples to analyze. You can filter the list of available samples by date, node, and data class. In addition, you can view information for each sample when you select it from the box below the displayed list. Information includes database path, UWS path, samples present in database, sample ID, and node sample ID (see figure 4).



Figure 2. The Retrieve window of the PMC interface

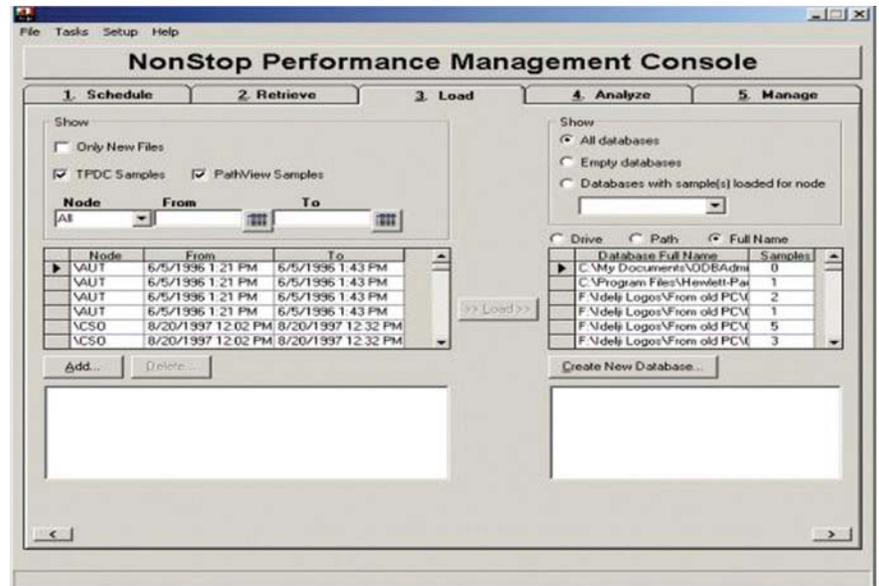


Figure 3. The Load window of the PMC interface

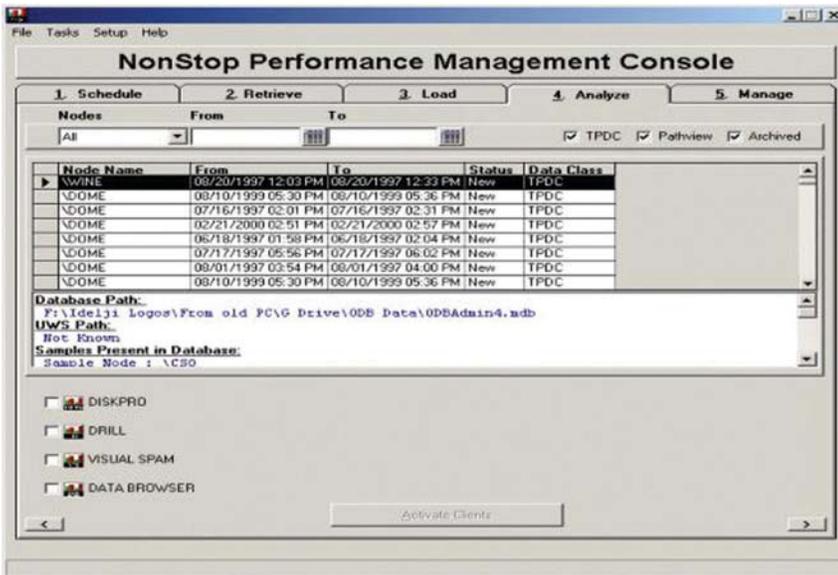


Figure 4. The Analyze window of the PMC interface

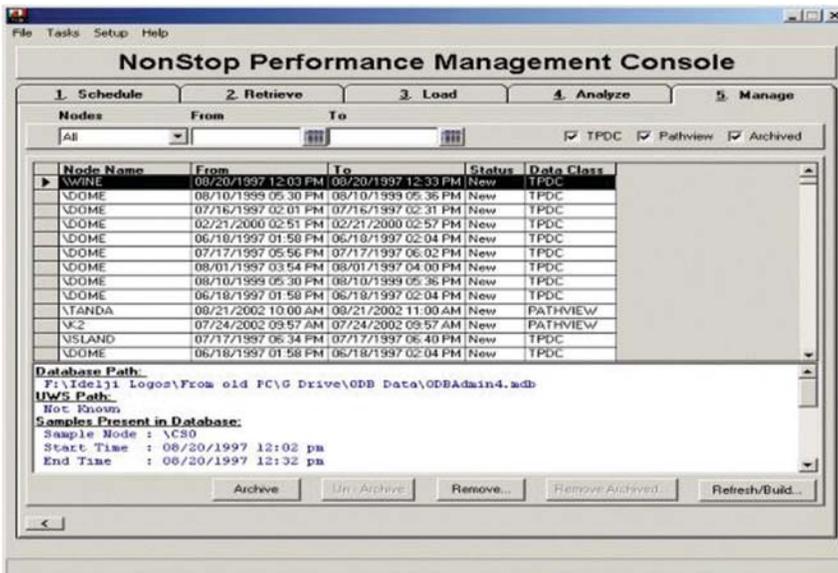


Figure 5. The Manage window of the PMC interface

managing samples and files

The Manage window of PMC allows you to manage the PMC-generated samples and ODB files on your workstation. The screen is similar to that of the Analyze function. It allows you to specify the type of samples you want to view. From this window, you also can mark or unmark specific samples as archived. You can remove data for one or more selected samples, or for all archived samples. You can even do a refresh or build, which searches your entire PC for all ODB and UWS data files and rebuilds the internal tables of PMC (see figure 5).

ordering information

product ID	description
SJ47V4	Open Database software

specifications

system requirements

Hardware	PC with Intel Pentium processor (or later version) with 128 megabytes of RAM and a 32-megabyte disk
Software	Windows Me, 2000, or XP, or Windows NT 4.0 operating system



For more information, go to www.hp.com/go/nonstop.

November 2002. Microsoft, Windows, and Windows NT are U.S. registered trademarks of Microsoft Corporation. Intel and Pentium are U.S. registered trademarks of Intel Corporation. All other product names mentioned herein may be trademarks of their respective companies. HP shall not be liable for technical or editorial errors or omissions contained herein. The information is subject to change without notice. The warranties for HP products are set forth in the express limited warranty statements accompanying such products. Nothing herein should be construed as constituting an additional warranty.

5981-4551EN

©2002 Hewlett-Packard Development Company, L.P.