
Introduction to IT Service Vision

<i>Using This Book</i>	1
<i>IT Service Vision Overview</i>	3
<i>How Can IT Service Vision Help You Manage Your IT Resources?</i>	4
<i>What Is a Performance Data Warehouse?</i>	5
<i>Tables, Variables, and Definitions in a PDB</i>	6
<i>Summarizing Data into Libraries in the PDB</i>	7
<i>How Long Is Data Stored in the PDB?</i>	9
<i>IT Service Vision Libraries</i>	10
<i>The Program Library</i>	10
<i>The CPSYSLIB Library</i>	10
<i>The Site Library</i>	10
<i>The SASUSER Library</i>	11
<i>The Performance Data Warehouse (PDB)</i>	11
<i>Additional Libraries</i>	12
<i>The Data Collection and Analysis Process</i>	13

Using This Book

Getting Started with IT Service Vision provides an overview of how you will use this product to manage and analyze your IT services. This book supports IT Service Vision software, as of Release 2.5.

Chapter 1 — Introduction to IT Service Vision

This chapter provides general information about this book and concepts about IT Service Vision that will be helpful to new and experienced users. It also presents terms and concepts with which all users should become familiar.

Chapter 2 — Starting IT Service Vision

This chapter provides information on starting the IT Service Vision client and server, accessing data on the client and server, and using the window interface.

Chapter 3 — Working with Your Performance Data

This chapter provides information for the IT Service Vision administrator who is responsible for gathering the raw data, creating the PDB, and processing and reducing data into the PDB. In addition, this chapter provides step-by-step instructions for completing the QuickStart, Create PDB, and Process/Reduce Wizards.

Chapter 4 — Analyzing Data through Reports

This chapter introduces you to various reporting tools that you can use to analyze your data. It also discusses the report styles that are available within these tools.

This chapter is beneficial for all users, but is particularly helpful for analysts who will be creating and running reports on the IT Service Vision client.

Chapter 5 — Additional Documentation and Other Resources

This chapter provides a list of additional documentation and other resources that you may find helpful while using IT Service Vision.

Appendix 1 — How to Access Online Information

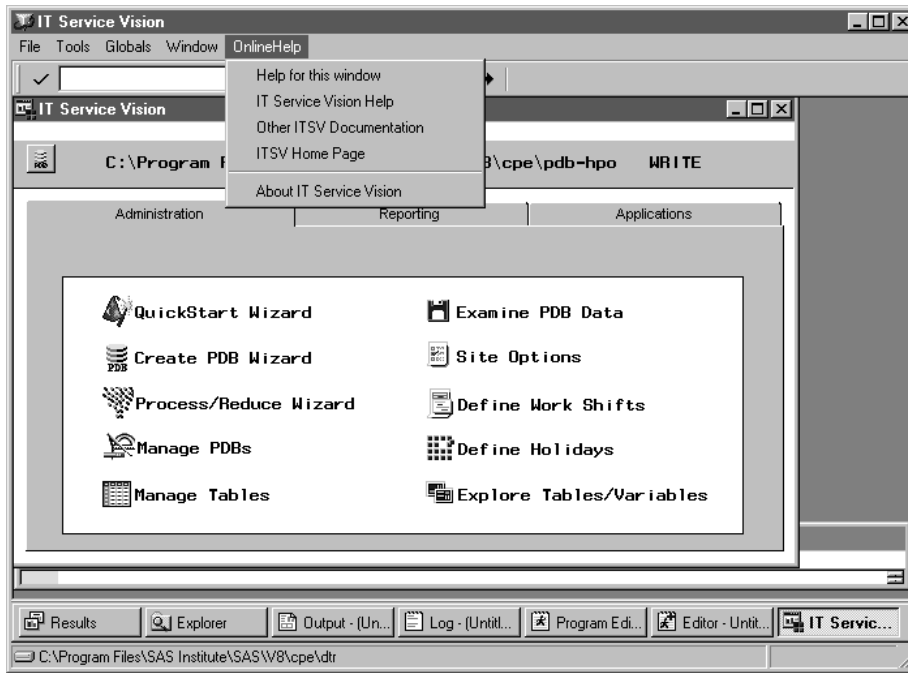
This appendix provides detailed instructions that enable you to access online information. Several types of online documentation are available including: IT Service Vision Help, Installation Instructions, Server Setup Guide, QuickStart Examples, and more.

Note: In this documentation, the terms MVS and OS/390 are used interchangeably - they both refer to the mainframe's operating environment. \triangle

To use IT Service Vision, you must also license base SAS software. This book is for IT Service Vision, Release 2.5 or later, running on SAS System Release 8.2 on a UNIX or PC host or on OS/390. IT Service Vision provides a GUI for UNIX and the PC and a GUI for OS/390. You can also view reports through your web browser.

Most of the sample windows shown in this book are based on the IT Service Vision GUI for PCs. Depending on your operating environment, your windows may differ from the samples shown in this book, especially if you are running on OS/390. Additionally, some of the tasks described in this book must be performed on a specific host, either the client or the server. This is described throughout the book, when applicable.

This book describes how to perform tasks using the window interface. However, most tasks that can be performed through the window interface can also be performed in batch using IT Service Vision batch macros. You can start the application in batch, process and reduce data in batch, and create and run reports in batch. The purpose and syntax for each macro are available online in the Macro Reference documentation. You can access this information by selecting **OnlineHelp** from the main IT Service Vision window:



Next, select **IT Service Vision Help**. The IT Service Vision Overview opens in the right panel of the window. Then, scroll down on this page and select **Using IT Service Vision Macros**.

For more detailed information about how to access online documentation about the macros, see Appendix 1, “How to Access Online Information,” on page 83.

IT Service Vision Overview

IT Service Vision is a performance evaluation tool that helps you to utilize your IT resources more effectively. It enables you to access, manage, and analyze large quantities of performance data. Performance evaluation is expanding to include areas such as troubleshooting, planning for growth, Web servers, applications support, and phone systems. IT Service Vision provides the technology that lets you summarize large volumes of data into a common format and then quickly and easily analyze the data using the reporting tools within IT Service Vision.

Using IT Service Vision, you can read and process data from almost any data source. IT Service Vision supports many popular network and systems management tools and data sources on OS/390, UNIX, and Microsoft Windows NT. It also allows you to read and analyze data from any time-based data source that contains a datetime stamp and whose log record formats are known.



The software that you use to gather observations about your IT services is referred to as a *collector*. IT Service Vision supports many popular data collection applications by providing tables for reading and storing the data. However, you can also create your own custom collector, also known as a generic collector, and IT Service Vision can read and process that data into your custom tables.

IT Service Vision reads raw data (metrics) in the format that they are logged by your collector or systems management tool. IT Service Vision can minimize the volume of data that you keep, by summarizing the detailed performance data into smaller groups or levels in the performance data warehouse (PDB). You select the statistics to be calculated at each summary level. After the data is reduced into the summary levels in the PDB and the statistics are calculated, you can analyze information over different time periods by running reports on a specific summarized group of your data.

You can also customize many aspects of the IT Service Vision application. You can use or modify tables and reports that are supplied with IT Service Vision and you can create new tables and reports.

How Can IT Service Vision Help You Manage Your IT Resources?

IT Service Vision can be used in a variety of systems management situations within your organization. You can use IT Service Vision as follows:

- to establish baselines for system use and performance. You can then track performance and workload volume against these values to anticipate the need for resource upgrades or to identify the causes of unsatisfactory service.

- to track activity and to highlight trends in usage and performance.
- to analyze your IT systems to determine the optimum resources required, so that you spend only the money and resources that are necessary to maintain excellent customer support.
- to run exception reports that can identify problems in your IT systems.

IT Service Vision features can help you to make effective business decisions concerning your IT services. A few of these features are as follows:

- You can process data from almost any data source.
- From a client, you can access and analyze data stored on a local or remote server.
- You can view reports through an IT Service Vision GUI or through your web browser.
- You can summarize data in the PDB, thus reducing your storage needs. You can then analyze and create reports based on the summarized data.
- You can point and click to access and navigate around your PDB and reports.
- You can work with demonstration PDBs for several popular data collection applications.
- You can create custom or drill-down reports with a variety of reporting tools.
- You can modify and create rules that identify exceptions in the data.
- You can access examples that set up various data sources for use with IT Service Vision.
- You can perform tasks in batch mode by using an extensive library of SAS macros. You can also perform tasks interactively through an IT Service Vision GUI.
- You can access the extensive analysis capability and support provided in the SAS System.

What Is a Performance Data Warehouse?

A performance data warehouse is a group of SAS libraries that contains the data used by IT Service Vision. (A performance data warehouse can also be called a performance database or PDB.)

The PDB contains data libraries, tables, and variables that define which statistics are calculated for the data. The PDB also stores options associated with the data.

Your site can have one or more PDBs, depending on how you want to collect and store your data. These are a few of the reasons why you may decide to employ multiple PDBs:

- You can use a separate PDB for each data source, such as Windows NT servers, UNIX servers, and phone switches.
- The data in the separate PDBs can each make use of different variables.
- Different groups may want to manage their own separate PDBs.
- You may want to create a test PDB so that you can investigate specific performance problems or facilitate testing.

Alternatively, you can choose to combine data from multiple data sources in the same PDB. You can organize your data and PDBs to meet your site's needs.

IT Service Vision supplies many table and variable definitions. These supplied definitions are the default specifications that govern how these tables or variables are created. When you create a PDB at your site, you can perform the following tasks:

- Add the supplied table and variable definitions and use them in your new PDB.
- Add the supplied table and variable definitions to your new PDB and modify them for your specific needs.
- Create new table and variable definitions in your PDB.

For example, if you add a supplied table and its supplied variables to your PDB, you may want to change defaults such as the age limits of data in each data level, the settings that specify which data is kept at detail level, and which statistics are calculated at each summary level. You may also want to delete variables that you do not want to use or to create new variables in order to calculate specific statistics.

The task of setting up your PDB is described further in “Using the QuickStart Wizard” on page 23.

Tables, Variables, and Definitions in a PDB

A **table** is a group of data, statistics, and definitions that is stored in a PDB. Each table contains data related to a specific performance area, such as a specific type of network activity. Each table also has a **table definition**, which is a description of the table’s properties.

Table definitions include:

- The name of the collector that is used to gather the data
- The type of table — interval or event
- The age limits of data in each data level
- The list of variables that are stored in the table.

Variable definitions are also associated with a table. A **variable definition** is a description of a variable’s properties. Variable definitions include:

- The data type, which defines whether the variable is numeric or character
- The interpretation type, which is used to determine the default statistics kept for the variable and how the variable is calculated in the reduction levels
- The format, for external display
- The length, for internal storage
- The statistics to be calculated for the variable at each summary level in the PDB.

You can use the GUI or batch mode to add supplied table definitions to your PDB. You can also modify table definitions and create new ones interactively or in batch. The table and the variables within the table are created automatically by IT Service Vision from the definitions.

You can also create definitions for formula variables and derived variables.

Formula variables are variables whose values are calculated from the data that you collect. These variables are calculated dynamically when you access the data so that you do not need to store the values of these variables in the PDB. **Derived variables** are also variables whose values are calculated from the data that you collect. After they are computed, however, the values of the derived variables are

stored with your data in the PDB. IT Service Vision does not have to calculate them again each time they are requested.

Within a table, data is stored in five libraries (also known as levels within a table): DETAIL, DAY, WEEK, MONTH, and YEAR. The data in the detail level is typically very similar to the data as it was recorded by your collector. The other four levels, also known as reduction or summary levels, contain your reduced or summarized data. The day level contains data summarized by the day, the year level contains data summarized by the year, and so on. For example, a table for process statistics may contain detailed information for each process in the detail level and also summarized data in each of the four summary levels (day, week, month, year).

When you display the data from one level of a table in the PDB, this “picture” of the data is referred to as a *view*. A view is named using a combination of the level and the table name in the form of *level.table*. For example, the view for the detail level in a table named NTPRCS would be DETAIL.NTPRCS. The view for the day level of the same table would be DAY.NTPRCS. When you report on your data, the reports are based on the views.

On UNIX and PCs, you can see a list of tables in a specific PDB, as well as status information about the tables, by selecting **Manage Tables** from the Administration tab in the main window of IT Service Vision. To see a list of variables for one of the tables, select a table from the list and select this menu path:

Locals ► List Variables

On OS/390, you can see the corresponding information by selecting

PDB Admin ► Config Active PDB Dictionary

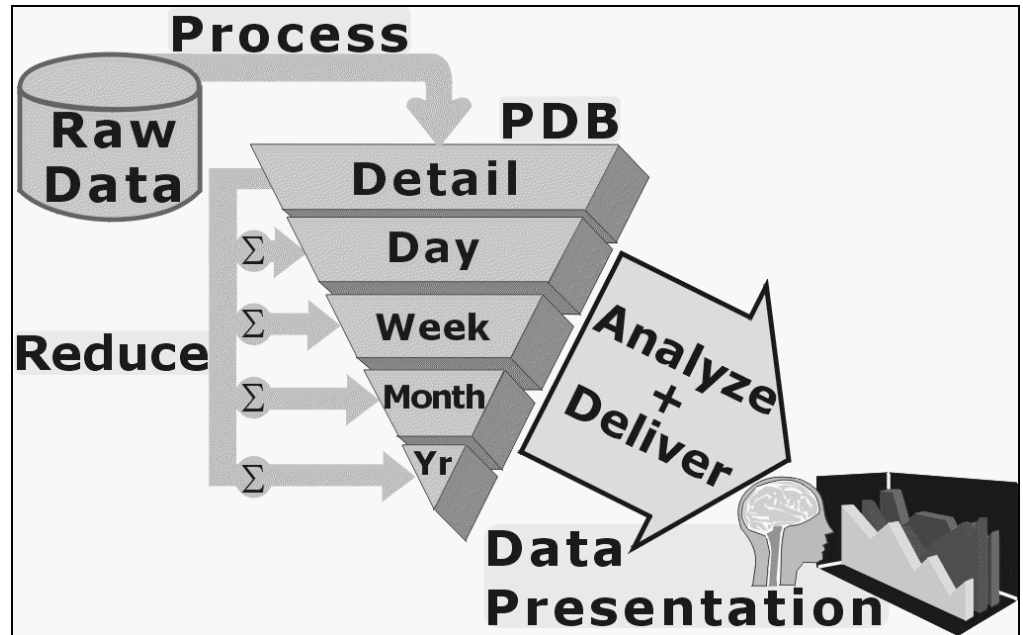
Next, select a table from the list, and then select

ItemActions ► List Variables

Summarizing Data into Libraries in the PDB

As previously mentioned, data is stored physically in five data libraries within tables. These libraries are known logically as *levels*. When your data is first transferred into the PDB from your raw data file, this is known as processing your data into the PDB, or the *process* task. This processed data is stored in the detail level library. The data in the detail level is very similar to the original data that you collected, with simple transformations (when applicable), such as converting continuously ascending counters into rates per second or adjusting a counter that has exceeded its maximum and restarted at zero. For collectors with vast amounts of data, such as transaction records, the process task may also perform some pre-summarization.

After the data has been processed into the detail level of the PDB, you then summarize or *reduce* the data into the other levels of the PDB. These other four levels are called *summary* or *reduction* levels, and they are day, week, month, and year.



The DETAIL library

contains your data after you process it into the PDB. Each observation in a data set represents an event or an interval. The data in the detail level is very similar to your data records in the raw data file.

The DAY library

contains daily statistics summarized by grouping variables, such as machine, hour, and day. These statistics are calculated from the data in the detail level. For example, each observation in a data set in this library may contain statistics for a specific machine, and for a specific hour of a specific day.

The WEEK library

contains weekly statistics summarized by grouping variables. The statistics are calculated from the table's data in the detail level of the PDB. For example, each observation in a data set in this library may contain the statistics for a specific machine, for a specific hour of the day, for a specific shift, and for a specific week.

The MONTH library

contains monthly statistics summarized by grouping variables. The statistics are calculated from the table's data in the detail level. For example, each observation in a data set in this library may contain the statistics for a specific machine, for a specific hour of the day, and for a specific month.

The YEAR library

contains yearly statistics summarized by grouping variables. The statistics are calculated from the table's data in the detail level. For example, each observation in a data set in this library may contain the statistics for a specific machine, for a specific hour of the day, and for a specific year.

When data is “reduced” or summarized, it is copied from the detail level and reduced into the summary levels in the PDB. (These summary levels are independent of each other.) The data is reduced based on the statistics that you select for each summary level. You can use the default statistics, or you can customize the list of statistics that you want to be calculated for your data at each reduction level. When the data is reduced, the specified statistics are calculated, and these summary statistics (not each observation from your data) are stored at each reduction level.

Note: You do not have to use all levels of a table. Furthermore, depending on your requirements, you can have different variables and statistics at each level. △

Note: Be sure to make a backup copy of your PDB before you make changes to it, and again after you have customized it. △

How Long Is Data Stored in the PDB?

The detail library might easily become very large because this library contains the data as it is logged by your performance management tool. As you process new data into the PDB and reduce the data into the summary levels, you will want to decide how long to keep data at each level in the PDB.

To reduce the space required for the detail library and the other libraries, your site can set age limits on the data at each level. After data is reduced from the detail level into the reduction levels, you may not need to keep the detail level data, or you may want to keep only a few days of data at that level. Similarly, for the other summary levels, you can choose a time limit, such as 3 days, 10 days, a month, and so on. The data will be deleted from each level when its age limit is reached. Using each of the summary levels, you can still analyze and report on long-term trends for your IT systems, but you do so using summarized data instead of large quantities of detail data.

The following built-in age limits are assigned by default to the levels in the table:

- Detail — 10 days
- Day — 45 days
- Week — 15 weeks
- Month — 18 months
- Year — 5 years

You can change the age limits by following this path from the **Administration** tab within the GUI for UNIX and PCs:

Manage Tables ► *select a table* ► File ► Properties

From the Table Properties window, select the **Table Status** tab to review the age limit settings.

The corresponding path in the GUI for OS/390 is

PDB Admin ► Config Active PDB Dictionary ► *select a table* ► ItemActions
 ► Edit Definition

IT Service Vision Libraries

IT Service Vision uses several libraries that contain items such as software, data, user options, site options, and more. These libraries are described below.

The Program Library

The IT Service Vision program library (PGMLIB) is a SAS library that contains the IT Service Vision software for your current host. It contains the master data dictionary (supplied table and variable definitions), supplied report definitions, supplied rule definitions, and the IT Service Vision client or server software programs for your current host. Supplied table, variable, report, and rule definitions are provided for many popular data sources to assist you in using those collectors and data sources with IT Service Vision.

When you create a new PDB, it contains many default settings. Your site can customize those default options, customize the supplied definitions, create new definitions, and add table and variable definitions to your PDB from the master data dictionary (in PGMLIB). You can also install user-generated table definitions in the master data dictionary, and your site can easily add those definitions to other PDBs at your site.

When IT Service Vision is updated, the information in PGMLIB is updated. You cannot modify PGMLIB. However, you can access PGMLIB to view, search, or copy from the list of supplied table definitions, variable definitions, report definitions, and rule definitions.

Note: Be sure to make a backup copy of your PDB before you make changes to it, and again after you have customized it. \triangle

The CPSYSLIB Library

The CPSYSLIB Library is a SAS library that is used on the PC and UNIX platforms. The library contains the IT Service Vision UNIX/PC GUI programs, screen definitions, and client-related information.

It is not available on OS/390.

The Site Library

The Site Library contains your site-related information, such as site holidays, default work shift schedules, default start of week setting, and more. After you or others at your site create a new PDB, the default properties for that PDB are copied from the site library. When you create a new PDB, you can customize the settings for that PDB.

A default site library is created when you install IT Service Vision. Therefore, a site library is installed on your client hosts and a site library is installed on your server host. If you are processing data for multiple sites or groups within your organization, then you may have a separate SITELIB library for each site or group. Thus, you can customize the default settings for each group.

In the GUI for UNIX and PCs, you can view many of your site options from the **Administration** tab in the main IT Service Vision window. To view the name of the currently defined site library, select **Site Options** from the **Administration** tab. Select **Change Sitelib** to display the name of the current site library. You can modify the current site library from this window.

The corresponding path in the GUI for OS/390 is

PDB Admin \blacktriangleright Define Work Shifts

If your organization wants to change the defaults in the site library to customize one or more options, you should copy the default site library and make changes to the copy. To modify settings in your site library, you must have write access permission for this library. However, you can read and copy items from the active site library if you have read access to the library. For more information on accessing your site library or setting site-specific preferences such as holidays and shifts, refer to the online help topic, *Work with an IT Service Vision Site Library*.

Note: Be sure to make a backup copy of your SITELIB before you make changes to it, and again after you have customized it. Δ

The SASUSER Library

The SASUSER library is a SAS data library that is used by your interactive SAS session for storing user-specific information. Examples of this information are the locations of PDBs that you use and any customizations that you specify for SAS options.

When you start the SAS System, a SASUSER library is automatically created for you, if one does not already exist. When you customize keys, menus, or other items within the application, these modifications are automatically saved in your SASUSER library. By default, any source, log, or report files that you save are saved in this library.

Note: Be sure to make a backup copy of your SASUSER library before you make changes to it, and again after you have customized it. Δ

By default, this library is not used in batch mode on OS/390. For platform-specific information on the SASUSER library, refer to the SAS Companion documentation for your platform and your current version of SAS.

The Performance Data Warehouse (PDB)

In addition to the data libraries described in “What Is a Performance Data Warehouse?” on page 5, the PDB contains several other libraries related to the data in a specific PDB. These libraries must be present in a PDB, even if they are empty.

The ADMIN library

is available to all PDB users. Therefore, it may be a convenient place to store graph and text reports, report palettes, logs, and so on, within a PDB. If you are using the data filtering routines to prevent duplicate data when you process your data into the PDB, this library is used by those routines.

The COLLECT library

is used for temporary data storage while your data is being processed into the PDB. For example, you can use the COLLECT library to store any staging data sets that you create for use by the Generic Collector Facility software. This library may also contain information or libraries used with MXG software.

The PDBWORK library

is used by the reduction step for intermediate data. The PDBWORK library (which is permanent) is used instead of the WORK library (which is temporary) so that reduction can be restarted from a checkpoint if necessary. Sometimes PDBWORK contains temporary SAS data sets, which are used for calculating statistics for the day, week, month, and year levels.

The data dictionary

is in each PDB. Also known as the DICTLIB library, the PDB's data dictionary contains general PDB options as well as information about the data that you have stored in each level of the PDB. For example, the DICTLIB library contains specifications about which data to keep in the PDB, and a count of the number of observations in each level of the PDB.

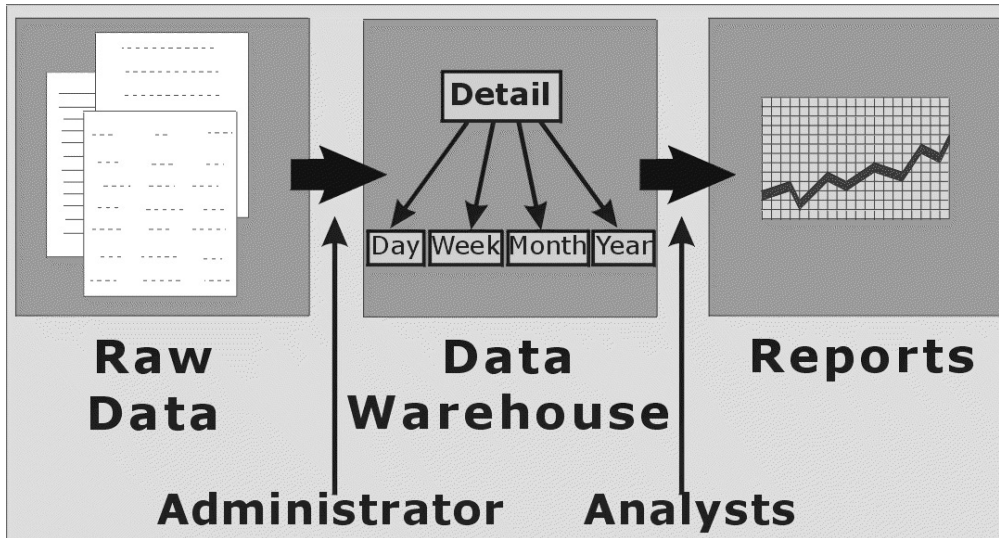
Additional Libraries

Depending on your host, additional items may be stored in other libraries.

- On UNIX and Windows NT hosts, additional software may be stored in an auxiliary library.
- Items such as MXG formats may be stored in an auxiliary library.
- If you have more than one release level of IT Service Vision installed, you will have a separate version of the PGMLIB library for each release.
- If you elect to archive data, archive libraries are created.

- When you run the QuickStart Wizard, the Wizard creates additional storage locations for materials, such as programs.

The Data Collection and Analysis Process



The routine tasks that you will perform to process and reduce data and view reports can be summarized as follows:

- 1 *Set up your data collector to work with IT Service Vision.* You can use the QuickStart Wizard to create your PDB, configure your site's PDB options within IT Service Vision, and create batch jobs that can be used to process data, reduce data, and create reports based on your data.

If you use the QuickStart Wizard, the process, reduce, and even the reporting tasks mentioned in the following steps can be performed through daily jobs that are created when you run the QuickStart Wizard. Based on the data category you select in the Wizard, the Wizard also provides an instruction file to assist you in setting up your collector with IT Service Vision. For information on using the QuickStart Wizard, see "Using the QuickStart Wizard" on page 23.

If you do not use the QuickStart Wizard, information on setting up your data collector is available in the *IT Service Vision Setup Guide* (see Chapter 5, "Additional Documentation and Other Resources," on page 77).
- 2 *Collect the IT systems data that you want to analyze.* You can collect data using an existing performance management tool, or using a custom tool that you write for a specific data source.
- 3 *Process raw data into the detail level of the tables in your PDB.* Your site can create a daily job to read and process the data, or you can use the daily job that is created when you run the QuickStart Wizard.

- 4 *Reduce the detail level data* into one or more summary levels in the PDB: day, week, month, or year. In most cases, you or an administrator at your site will use a job generated by the QuickStart Wizard, or will set up a daily job to reduce the data immediately after it is processed into the PDB.
- 5 *Run supplied report definitions or create and run your own custom report definitions* to help you analyze your data and identify performance issues. You can generate reports using data from any level of the PDB.