About the Precise PMDB

Analyzing resource consumption over a long period of time is the only way to predict the future resource consumption and response time of your environment. To identify resource consumption trends and patterns, you need to see summarized data on resource consumers in your environment over various periods of time.

The PMDB helps you identify resource consumption patterns and predict future resource consumption and response times by:

- Tracking historical resource consumption trends so you can understand and predict long-term performance behavior.
- Performing period-to-period comparisons so you can analyze performance improvements or performance degradation over time.
- Tracking load patterns, entity changes, entity statistics, and component parameter changes so you can understand their effects on performance.
- Tracking data growth and data distribution changes so you can optimize data storage management.
- Proactively detecting performance bottlenecks before they turn into problems and issue alerts when performance degrades from established

By using the PMDB to collect this information, you can manage your applications better, make knowledgeable decisions about application changes and hardware upgrades, and improve plans for the future.

The PMDB collects and manages the PMDB data by:

- Scheduling and running batch processes that load historical data from each instance on each application tier (AppTier)
- Managing the summary procedure
- Providing a common interface for requests coming from Precise user interfaces
- Managing and maintaining the PMDB database by filtering unnecessary data, saving baseline data, summarizing data, and purging old data

You can control the settings of the PMDB from AdminPoint. AdminPoint is the central administration console of the Precise product suite. For information about AdminPoint, see the *Precise Administration Guide*.

Instances in the PMDB

The PMDB identifies statistics data by the instance from which it is collected, such as an Oracle database or a SQL Server installation. Before you can load data using the Data Loader, you need to define the instance to which the data belongs. For information on how to define and view instances, see Defining instances in the PMDB.

Once you have defined an instance, you can view the instance ID. You must provide this ID when loading information into statistics tables (see Columns Required for Statistics Tables).