The Sybase/IBM Data Mall

By Ron Wastal

The Sybase/IBM Data Mall architecture gives organizations the ability to deliver rapid, new business value and at the same time build collaborative and effective data management and movement strategies. It delivers fast and cost-effective interactive "store fronts" with a high-quality and manageable infrastructure "warehouse."

Data warehousing implementations are only as good as the quality and timely information they deliver to business users. Balancing the corporation's need for centralized coordination and control with each department's need for distributed empowerment and execution can be difficult.

With Sybase's parallel data warehouse database (Sybase MPP™) on the IBM RS/6000™ SP, along with Sybase's new interactive data mart database (Sybase IQ™) on the new RS/6000 Symmetric Multiprocessor (SMP) Server, combined with leading edge middleware products (Sybase® Enterprise CONNECT™), users and Information Technology (IT) professionals can now get their data warehouse "cake" and "eat it too."

Rethinking the Data Warehouse Concept

Data warehousing is an acknowledged technology in the corporate arena, accepted by customers, consultants, and vendors as a unique tool for achieving competitive advantage in the global marketplace. However, delivering on the promises of data warehousing, including more effective target marketing, more precise financial forecasting and budgeting, more efficient channel and logistics management, and more intimate customer service, can be a daunting challenge for IT executives.

A data warehouse was originally defined as a subject-oriented, integrated, time-variant, non-volatile collection of data that is specifically collected to aid management decision-making processes. The promise was to deliver information on how to discover or respond to a product opportunity quickly, a competitive threat or weakness immediately, or understand new opportunities and marketplace trends very early in the cycle. This was done by mining enterprise data stored by various departments within an organization. The principle of the data warehouse was to put all this data into a common repository—the data warehouse—thereby making it more valuable as a business information resource versus a departmental operational resource.

Unfortunately, while the IT group set out to build a common resource accessible to thousands of users with a wide variety of skill levels and interests, the business users (under immense business pressure) did not wait, deciding instead to implement their own departmental data warehouses, later termed "data marts."

Data Warehouse Evolution

First generation data warehouses were usually developed around centralized, mainframe-based architectures that supported both transactional production systems and decision support. This architecture proved to be unworkable from a practical perspective since the data models defined in the warehouse were often very complex. In addition, with Decision Support Systems (DSS) and Online Transaction Processing (OLTP) requests competing for the same computing resources, the systems were unable to deliver satisfactory performance.

Second generation data warehouse systems evolved when IT departments consolidated data onto a single mainframe platform or commercial parallel processing systems. At the same time, independent departments or divisions pursued their own mini-warehouses (data marts) to support their analytical needs.

Data marts were especially attractive to end users since they were often easier to fund, faster to implement, and focused on delivering specific business value to a targeted set of knowledge workers. Information in a data mart is consolidated from production-level information organized to directly reflect the departmental understanding of business process, workflow, and so on. Data marts can be built rapidly but, over time, must be implemented to share common data elements, transformations, and management tools.

Now, a third generation of data warehouses is being constructed, with emphasis on creating a cooperative environment between the corporate data warehouse and the individual data marts. In this Data Mall approach, data marts are built rapidly but with an architecture that promotes sharing of common data elements and transformation programs. Over time, as multiple data marts go into production, the enterprise warehouse becomes an optional port of the architecture to hold large amounts of commonly accessed data.

This hybrid model focuses on delivering high performance, system scalability, and manageability at both the warehouse and data mart level, while also delivering highly responsive and flexible query performance to satisfy the expanding needs of end users. This last requirement is crucial since a resource like a data mart is difficult to model over time—users often do not know what they want to do with the data mart until the data is in their hands. This means that the interface tools and underlying architecture of the data mart must permit users to explore the data interactively, learning and discovering about the resource available to them.

Specifically, Sybase MPP running on the IBM RS/6000 SP provides a scalable architecture that supports the data intensive, high-capacity, high-volume batch and pre-defined query processing typical of very large data warehouses. Simultaneously, Sybase IQ running on a IBM RS/6000 or RS/6000 SMP Server delivers exceptionally fast query response in a flexible, ad hoc manner that addresses the demanding information analysis needs of the data mart user.

The Data Mall—The Next Generation

A *data mall* is a data warehouse/data mart environment combined under a single system roof. It uses the unique architecture of the IBM RS/6000 SP to prototype, test, implement, and if necessary, distribute these rapidly built data marts in a manageable fashion. IT managers can now optimize and tune their data mart structure and evaluate LAN/WAN traffic concerns before distributing systems. This assures end-user satisfaction and provides faster, more reliable implementations. With Sybase's optimized database middleware (Enterprise CONNECT) and server products coordinating the warehouse (Sybase MPP) and data mart (Sybase IQ), system administration requirements are greatly reduced and total system integration is simplified. The following sections explore this solution in greater detail.

A data warehouse decision support environment dictates performance requirements in five key system evaluation areas: scalability, query performance, capacity, availability, and administration.

IT managers can now optimize and tune their data mart structure and evaluate LAN/WAN traffic concerns before distributing systems.

Scalability

The continuing acceptance of decision support systems is fueling a rapid growth in database size. Surveys indicate that a significant percentage of companies have started or intend to implement a data warehouse system in the next few years. In addition, companies continue to store increasing volumes of production and transaction data as their data-capture facilities improve.

With current data warehouses managing hundreds of gigabytes to a terabyte or two, most industry experts predict that future warehouses will be forced to manage multiple terabytes of data. By logical progression, data marts must also be prepared to scale, handling additional capacity from the megabyte- and gigabyte-class loads they deliver today.

Query Performance

Data warehouses and data marts often need to support three broad classes of queries:

- **Interactive, ad hoc queries:** Read-only requests accompanied by expectations of near real-time response where an analyst explores data via a series of drill-down or drill-up queries, searching through various levels of data, usually covering narrow time frames or search categories.

- **Complex queries:** Processing requests that cover large amounts of data spanning broad time periods or product/service categories. These computationally intensive requests typically include join, sort, or group-by operators and are, in some cases, performed on a repetitive basis. In spite of the complexity, it is assumed that rapid turnaround is provided, since information value (accuracy) is related to the timeliness of the analysis.

- **High-speed, mixed workload queries:** Requests that require near real-time response and read/write support across multiple platforms. Typical applications are online customer support or order entry, where an operator will need to quickly search a customer record from a large list, then modify and enter data, or access a more complex data resource such as a Help Desk Q&A record.

These variations describe highly variant performance demands and often require different systems to provide effective support to diverse end-user requirements.

Sybase MPP is an extension of the Sybase SQL Server technology specifically designed and optimized for massively parallel processing environments.

Capacity

Both data warehouses and data marts are typically updated on a batch basis, using periodic, scheduled loading times. Routine operating requirements, such as backups, must be scheduled. Since the update process involves large data blocks, the information architecture must provide sufficient load/update capacity to handle the batch requests in the allotted time period.

Availability

Construction of a data warehouse and data mart environment requires a significant investment that only returns value when it is operating. While the system hardware/operating platform and database may not require total availability (fault tolerance), it should have high availability with failure recovery mechanisms that let users continue to operate even if a component of the system is down.

Administration

Maintaining data integrity and availability in a distributed environment can be difficult. For example, data marts should be continually re-evaluated for their effectiveness; that is, can the data mart support greater than 95% of user queries, or do some routine drill-through (back to the warehouse) or drill-across (to other data marts) requests place excessive demands on LAN/WAN networks or high-speed interconnect capacity (in the data mall), and burden the data warehouse with ad hoc requests?

Sybase and IBM

Sybase and IBM offer a unique suite of complementary products that deliver the performance and availability that users demand.

Sybase MPP and IBM RS/6000 SP

Managing the capacity demands of a data warehouse while offering superior data mining and analysis capabilities on an open platform can seem like a big challenge. Traditional parallel database systems use multiple processors to access data using either shared memory or shared disk architectures. This approach, however, limits performance, especially as the database scales in size. In fact, the shared-nothing architectures of current generation massively parallel hardware can support data sizes and throughput rates an order of magnitude greater than shared memory and shared disk databases. Because of software limitations, massively parallel hardware for database applications has been limited to niche applications.

The RS/6000 SP represents a next-generation hardware design that uses widely available components and operating systems to provide massive processing power, enormous scalability, and exceptional availability in a highly affordable package. This makes parallel database architectures affordable in massively parallel environments. However, parallel databases do not always take full advantage of the RS/6000 SP architecture.

Sybase MPP is an extension of the Sybase SQL Server™ technology specifically designed and optimized for massively parallel processing environments. The open, portable, parallel architecture of Sybase MPP delivers unmatched performance, unlimited scalability, and database management utilities to support Very Large Database (VLDB) operations. Sybase MPP uses multiple SQL Servers, each accessing one or more processors, to implement a parallel shared-nothing database environment.

Performance

Sybase MPP delivers superior performance for pre-defined and batch queries.

- Complete, parallel, shared-nothing architecture takes full advantage of the massively parallel

design to deliver exceptional processing performance by minimizing system bottlenecks such as memory or I/O contention.

- Integrated system utilities simplify and speed system tuning for optimum use of hardware resources.

- Fully parallel operations—from utilities and queries to data operations such as insert, update, or delete—deliver flexible performance to support end users while handling real-time systems management and data updates.

- Parallel loading (parallel bcp) reduces load times on bulk updates and assures timely availability for user queries.

Scalability

Sybase MPP scalability makes it ideal for growing data warehouse facilities.

- Design scales to near linear levels in both response time (speed up) and system throughput (scale up) as additional processors, workload, and data are added to the system.

- Parallel performance enhancements are transparent to users, protecting investments in applications and training.

- Sybase Open Client/Open Server architecture supports hundreds of existing end-user analysis tools.

VLDB Management Utilities

Integrated management utilities simplify delivery of parallel performance using Sybase MPP Manager and Configuration Services.

- Configuration Services helps customers analyze workload requirements and recommends optimal software and hardware configuration.

- Sybase MPP Manager provides graphical utilities to help administer the system.

- MPP incorporates tools for increased availability, including automatic failure recovery from disk, process, or system node errors (hierarchical monitoring).

The RS/6000 SP is a general-purpose, scalable parallel system that can incorporate between 2 and 512 nodes. Each RS/6000 SP node—packaged up to 16 per frame—is a self-sufficient system, containing local memory and disk capacity, and using the latest PowerPC™ or POWER2 processors. Each node runs its own copy of AIX® and other RS/6000 system software.

The heart of the IBM RS/6000 SP is the optional SP switch, a high-bandwidth, low-latency switching network that binds the nodes together. The switch incorporates a unique combination of topology and architectural features to scale aggregate bandwidth as the system grows. As nodes are added to an SP, switch performance scales at a near linear rate, which delivers peak node-to-node bandwidth for TCP/IP traffic of 70 MB per second, and bandwidth in excess of 100 MB per second for Message Passing Interface (MPI) protocol.

On a 16-node SP system, this equates to an aggregate bi-sectional bandwidth of over 800 MB per second. The SP switch design has point-to-point interprocessor communication time independent of the relative location of the communicating node. The resulting near linear scalability is unmatched by direct networks such as rings and meshes.

The RS/6000 SP's performance-oriented architecture is designed to support large-scale data warehouses.

Optimized for Data Warehouse Applications

The RS/6000 SP's performance-oriented architecture is designed to support large-scale data warehouses. On-chip and off-chip caches compensate for the highly mixed workloads and data sets typical of RDBMS applications. This ensures that processor cycles are used for database operations, not for retrieving frequently used instructions and data. Asynchronous I/O minimizes waiting for disk or network availability.

To optimize decision support, AIX has been designed to reduce lock overhead, allowing greater simultaneous system access. A dynamic kernel eliminates the need to take servers off-line for many routine configuration changes. The Logical Volume Manager (LVM) enables simplified filesystem expansion, data distribution, and mirroring.

In addition, the RS/6000 SP delivers unsurpassed availability by extending the benefits of conventional High Availability Cluster Multiprocessing (HACMP) architecture, including automatic workload balancing in the event of a node failure, to its full architecture. Instead of the standard 8-way limitations of HACMP structures, the RS/6000 SP can extend this to its full 128-node (or higher) architecture.

Data Marts on Sybase IQ

While parallel database technology is ideally suited to managing large capacity and complex queries, this database architecture often has difficulty scaling concurrent, mixed-user workloads. Flooding a parallel system with continuous interactive queries can drain resources and slow response time on all processing functions. The logical evolution is to pump data from the main data warehouse into smaller databases (data marts), designed specifically to support the needs of discrete end-user constituencies. These individual data marts can then be distributed to handle ad hoc queries of individual knowledge workers.

Sybase's new interactive database, Sybase IQ, is a tool for managing interactive queries in a high-volume, rapid response environment. IQ's technology reduces the amount of work per query required by the system, providing query performance improvements 5 to 500 times faster than traditional RDBMS systems. Plus, with Sybase IQ, the system can be tuned once, versus tuning for each query; this simplifies setup and administration. Coupled with the IBM RS/6000 SMP Node, Sybase IQ delivers rapid response at a fraction of the cost of stand-alone data warehousing solutions.

Sybase IQ can scale the data mart to support from tens to thousands of users on a standard SMP platform.

The key to the breakthrough performance in Sybase IQ is Bit-Wise™ query processing. A step beyond traditional indexes or bit-map indexes, Bit-Wise technology represents all values as bits in indexes. A comprehensive five-part indexing scheme manages various data types, optimizing the indexing scheme against cardinality. Indexes can be applied concurrently versus traditional relational databases that are generally limited to one index per query.

Sybase IQ's high cardinality, bit-wise processing approach provides rapid, dynamic aggregation of relational data, delivering vastly improved query response. Column-wide data storage and compression also provides up to 98% reduction in disk I/O for fast response while eliminating table scans and unpredictable results. Hot column caching supports dramatic performance gains on iterative, drill-down/drill-up analyses. Sybase IQ's open architecture is also compatible with a broad range of front-end query tools.

The Scalable Data Mart

Sybase IQ has the unique ability to scale the data mart to support from tens to thousands of users on a standard SMP platform, such as IBM's RS/6000 SMP Server. By eliminating table scans and making decision support a CPU-bound rather than an I/O bound process, Sybase IQ allows the data mart to deliver predictable support for many more users.

The software also provides simplicity of management and flexibility. Data mart configuration designs are often based on perceived user requirements. However, once users have access to the data, their need to access data beyond the originally designed data mart becomes evident. Excessive drill-through or drill-across requests indicate non-optimized mart construction.

IQ simplifies reconfiguration by supporting heterogeneous data sources and delivers administrative tools for system configuration and optimization. Since tuning is data dependent, not query dependent, re-tuning of the system is required only when the data definitions are changed.

The Data Mall and Beyond

The scalable RS/6000 family enables the warehouse and mart environments to be implemented on a common, binary-compatible platform. This means you can grow a system from either direction: start with a data mart or two then upscale to a data warehouse, or start with a data warehouse and build/feed data marts.

The RS/6000 scales from uniprocessors to SMP to clustered processor configurations to massively parallel systems like the RS/6000 SP. Since the same architecture and operating system are used on all systems, application portability is enhanced, and system training and support requirements are reduced.

Some advantages of the RS/6000 SP go beyond performance. While a data mart is a powerful concept at the user level, maintaining synchronous data and assuring overall system performance presents a challenge. Using the RS/6000 SP, IT managers can prototype their data warehouse/data mart construction by creating a data mall—a data warehouse and data marts coordinated by middleware under a single system roof.

Figure 1 shows how most data warehouses and data marts are being implemented today—as competitors attempting to attract information shoppers. Data is delivered to both the data warehouse (driven by IT needs) and data marts (driven by users needs) from a variety of departmental operational systems. The data marts often receive the bulk of the "shopping" since they are designed around specific business needs, while the data warehouse struggles to compete (and fund its existence) due to the lack of user interest.

Figure 2 depicts an evolutionary scenario implemented on an RS/6000 SP, with multiple nodes dedicated to the core data warehouse. It is running a Sybase MPP parallel database with distributed RS/6000 servers running Sybase IQ to handle data mart queries from end-user workstations. This allows individual departments to have some level of autonomy and control over their local data mart and receive the dedicated query performance needed to

Current Implementation of a Data Warehouse

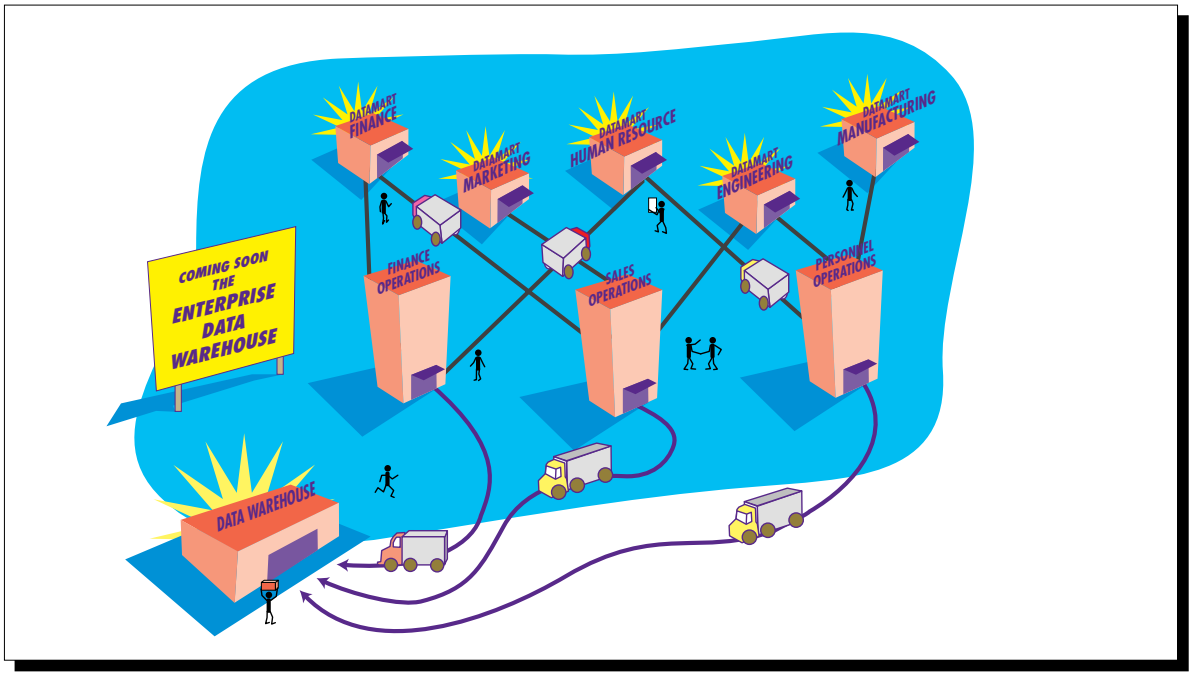

Figure 1. Current implementation of a data warehouse

Sybase MPP on RS/6000

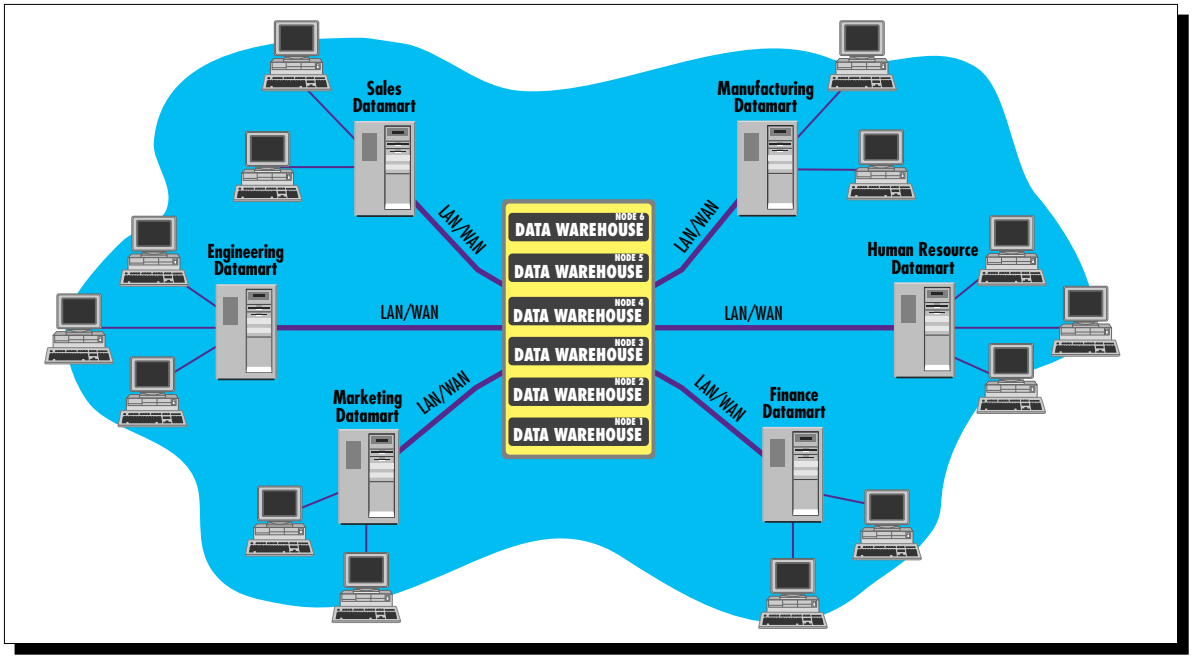

Figure 2. Sybase MPP on RS/6000 feeding RS/6000-based data marts running Sybase IQ

make the system valuable, all while maintaining centralized data warehouse quality and consistency.

Figure 3 introduces the latest evolutionary concept built around centralized coordination with distributed capabilities—the data mall. Imagine the impact of having a retail warehouse and individual storefronts all located under a single roof. Distribution is "hyper-accelerated" and simplified since distribution is just "down the hall" or over the interconnect. The data mall manager can clearly observe inventory traffic and evaluate demand at both the store and warehouse level.

The translation of this concept to an actual implementation on an RS/6000 shows how the individual data marts can be prototyped on an RS/6000 SP node. It is easy to clearly assess functional performance, data traffic requirements, and user demand before the system is considered for physical distribution to a remote field location. This permits proper optimization of the database design and preparation of the LAN/WAN for the new traffic patterns.

Eventually, distribution of the data mart might occur on a geographically remote stand-alone RS/6000. Since all RS/6000s are binary compatible, this transition is smooth, regardless of the prototype and distribution system configurations.

The concept of the data mall illustrates just one example of the potential that the combination of the RS/6000 SP with Sybase MPP and Sybase IQ offers those seeking to implement decision support strategies.

Essential Infrastructure in the Data Mall

Critical to the effectiveness of any data warehouse architecture is the controlled movement of data between individual data marts and the corporate data warehouse, or between individual data marts. Movement of data in any one of these architectures has the following characteristics:

- High volume, high performance
- Bulk load, unload
- Replication, update, refresh, enhance, expand, automated, and scheduled

Sybase InfoPump™ offers one way to accomplish data movement and materialization within a warehouse environment. InfoPump is a cross-platform tool (licensed from Platinum Technology®) that is used to automate the movement of bulk data on either a timed basis or through triggers that operate under defined criteria. InfoPump also has features that allow for the transformation of the data after it is captured in the source file and

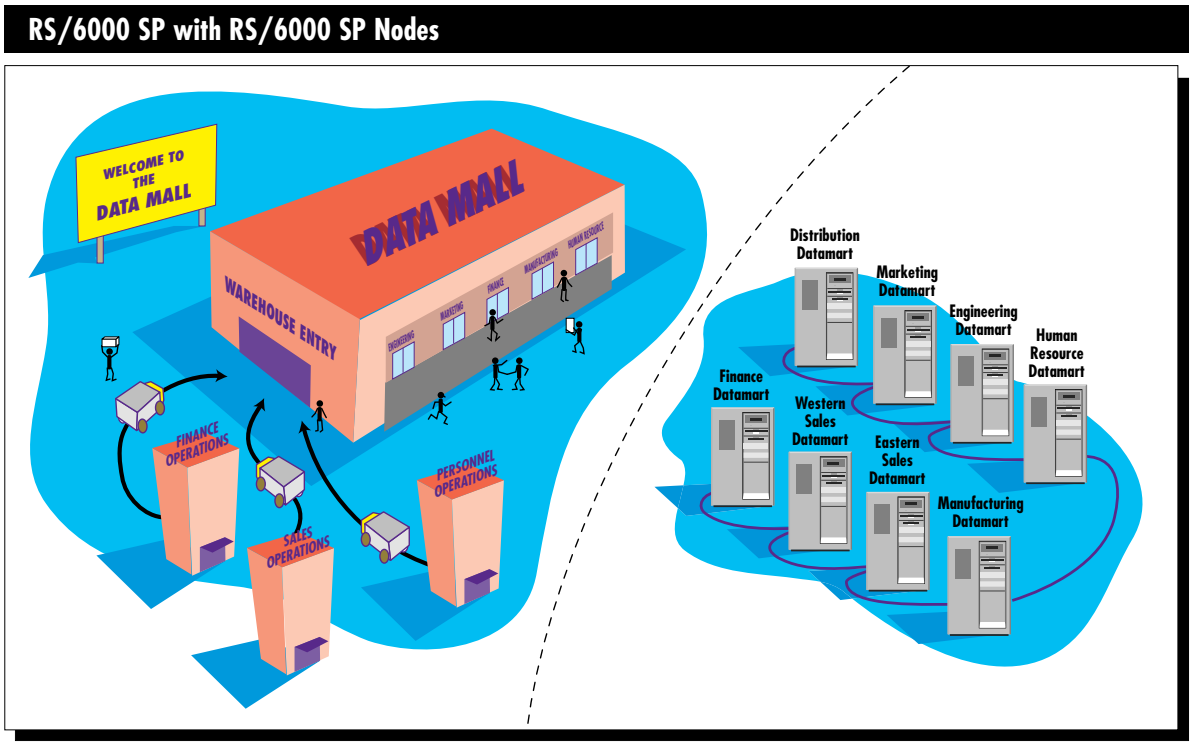

Figure 3. RS/6000 SP with RS/6000 SP nodes

Sybase on IBM Platforms

Sybase and IBM offer a complete set of solutions and services to address enterprise-class information services requirements:

Optimized Database Servers

- **Sybase IQ:** A powerful tool for data marts and interactive query support in a decision support environment; Sybase IQ runs on the IBM RS/6000 family and delivers rapid response for high-volume ad hoc queries; running on the IBM RS/6000 SMP systems, IQ is scalable to large numbers of users
- **Sybase MPP:** A "no-contention", parallel database designed for IBM's massively parallel processor system, the IBM RS/6000 SP; Sybase MPP excels at performing standardized queries across large data volumes and centralizes VLDB data warehouse functions
- **Sybase System 11™:** A family of database servers optimized for the IBM RS/6000 family and suited to supporting mixed-mode query and OLTP applications

Middleware Tools

- **Replication Server:** Replicates data across distributed systems for high availability, high performance, and data sharing across heterogeneous databases
- **InfoPump:** Performs summaries and aggregations to the data before delivery to data marts
- **OmniConnect™:** Read and update information from many heterogeneous data sources

Client Tools

- **PowerBuilder Products:** A complete family of tools for assimilating and presenting data, and developing applications for Windows™, OS/2®, UNIX®, Macintosh®, and Internet clients

before it is placed into the target database or data warehouse.

As a business environment continues to expand, data movement requirements shift, with increasingly specific requirements for systems and applications that have been installed. Huge volumes of operational data must be converted from its raw form and moved into the data warehouse. Companies are now looking to develop enterprise-level data warehousing solutions that integrate and manage a wide variety of tools and data, such as modeling and transformation, while simplifying the scheduled distribution of business data to the various data warehouses or data marts.

The Sybase/IBM Advantage

Sybase MPP and IQ, in combination with the IBM RS/6000 SP, deliver unmatched performance, unlimited scalability, high availability, and very large data warehouse management. These systems offer unique resources for prototyping, delivering, and distributing data marts while assuring that the systems will perform as users expect. This unique combination allows customers the ability to either evolve their information infrastructure top down or user up. Although top down seems like the more logical way to build a corporate architecture, user up appears to be a more pragmatic and lower-risk approach to delivering a comprehensive enterprise data warehouse.

Ron Wastal, Sybase, Inc., 6475 Christie Avenue, Emeryville, CA 94608. Mr. Wastal is director of Product Management for the Enterprise Database Products Group. He has been involved in over 50 detailed data warehouse projects with some of the largest corporations in the world.

VisualAge Tools for Java

VisualAge tools for the Java programming environment, available in late 1996, will provide a complete development environment, including an editor, debugger and browser, along with a Java class library and a VisualAge Data Access builder for visually constructing data access through the use of Java Database Connection (JDBC).

VisualAge software application development tools allow developers to construct applications visually by connecting pre-fabricated, reusable software components from an expansive library of predefined classes and parts from IBM and other vendors. The visual-construction-from-parts technology will be extended to the Java programming environment, enabling developers to visually build Java applets.