hp zero latency enterprise architecture

bridge the information gap

a white paper from hp

zero latency enterprise data store

contents

  - manageability of very large databases
  - mixed-workload support
  - robust integration of middleware components
  - ZLE demonstrations
  - complete life-cycle services
- conclusion

As companies struggle to survive in today's competitive environment, the winners will be those that understand their customers more intimately and that react faster and smarter than the competition. Your company's greatest assets in this battle are its information resources combined with its decision-making capabilities.

Companies today need information to support diverse business functions such as online business transaction processing; customer interfacing; and strategic, corporate, and departmental real-time decision-making. The question is, what information architecture best supports the decision-making capabilities of the enterprise? There are a number of choices, among them an enterprise data warehouse (EDW) combined with data marts or an operational data store (ODS).

Many companies have resorted to traditional data warehousing architectures (EDW and data marts) to solve their most pressing decision-making needs. Indeed, classical data warehousing does help an organization address many of the strategic, corporate, departmental, and knowledge discovery issues they face. But the problem with this architecture is that it does not enable what is perhaps the most critical enterprise capability today—real-time, or zero latency, decision-making. Applications in which a zero latency response to an event can provide a significant competitive advantage include attrition or churn management, billing, marketing campaign analysis and tuning, rescoring of customers, sales performance analysis, fraud detection, credit and risk management, and inventory tracking and management.

At the heart of a zero latency enterprise (ZLE) solution is the ZLE data store. This data store bridges the information gap between operational systems and the data warehouse. It contains data at the event detail level, coordinated across all relevant source systems and maintained in a current state. Thus the ZLE data store enables the integrated and up-to-date view of the business that is essential for a zero latency enterprise.

This white paper explores the role of the ZLE data store in the ZLE architecture and defines the types of informational needs it can satisfy. It explains how the ZLE data store fits into the corporate data architecture scheme and then examines the unique capabilities of the HP NonStop™ server to deliver on the ZLE data store promise.

the operational data store defined

In his book *Building the Operational Data Store,* Second Edition (New York: John Wiley & Sons, 1999), W.H Inmon defines an ODS as "an architectural construct that is subject oriented, integrated (that is, collectively integrated), volatile, current valued, and contains detailed corporate data." At first glance, that definition may seem cryptic, but upon closer inspection, it reveals the fundamental properties of an ODS.

- *Subject oriented*—Unlike operational data in most enterprises that is application specific, the data contained in an ODS is organized around subjects of enterprise such as customers, products, activities, policies, claims, and shipments.

- *Integrated*—Instead of "stovepiped" information specific to a single application, the ODS reflects a composite, integrated, cleansed, and enriched view of the entire enterprise.
- *Volatile and current valued*—Unlike a data warehouse, in which information is appended to the data store to reflect information as history, the ODS is constantly changing to reflect the current status of the subjects in the enterprise.
- *Detailed corporate data*—While the ODS can contain some aggregate or summary data, its real value lies in storing the *detailed* transactions of the enterprise. In fact, the definition of *detailed* is changing because companies are beginning to realize that they need to understand every customer interaction (for example, a mouse click on the website or a keypad depression on a wireless phone), regardless of whether or not it results in a completed transaction.

The ODS does not contain all information from all of the corporate data stores. Rather, it contains a subset of the enterprise information deemed essential to resolving key operational business issues. Because business is bound to change, the ODS must be designed such that new information can be added iteratively and painlessly.

Is the ZLE data store a conventional ODS? Absolutely not—the demands of a ZLE data store are significantly more complex and challenging.

ZLE requirements

A conventional operational data store or data warehouse is insufficient for a ZLE solution. It does not contain the current, detailed information necessary to conduct business with zero latency. Additionally, it does not provide for the integration of business processes and rules essential to inform all parts of the enterprise of current events and interactions.

A ZLE environment requires an architecture that merges enterprise application integration (EAI) with a ZLE data store so that the enterprise can

- Synchronize information across multiple enterprise applications
- Integrate data into a single storage cache for the enterprise with zero (or minimal) latency
- Provide data to business intelligence applications for analysis and data mining
- Apply knowledge generated from business intelligence analysis as real-time decision-making rules applied consistently across all channels of customer contact, products, and services
- Support massive online transaction processing (OLTP) access to up-to-the-second data that is integrated from multiple and disparate operational systems
- Route information intelligently between applications and their databases

In 2000, Compaq, now part of HP, introduced and implemented the first architectural model that supports ZLE requirements. As figure 1 shows, the ZLE data store, which serves as the memory and foundation of the HP ZLE framework, is tightly coupled with the EAI component. Without the ZLE data store, a ZLE solution cannot be created.

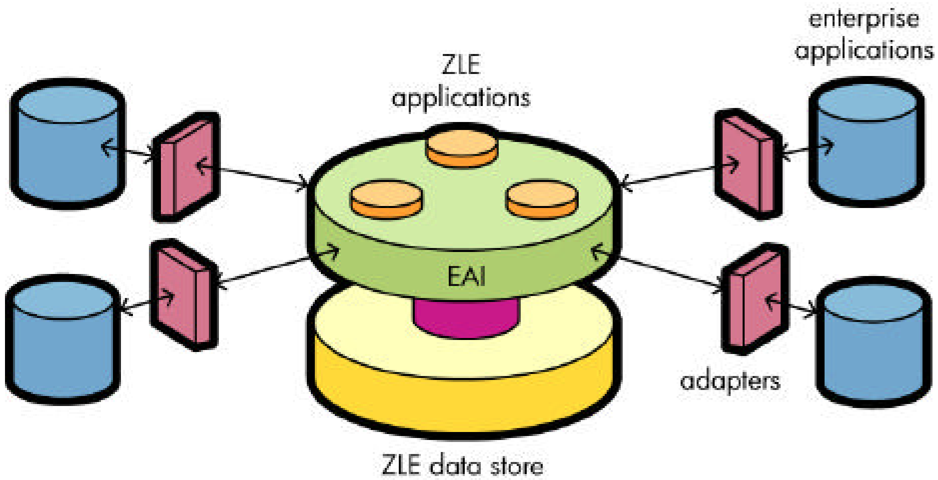

figure 1. the HP ZLE framework

A ZLE framework empowers companies not to simply follow their competition but to skip a generation and take a quantum leap toward a zero latency enterprise. It enables

- A single view of information from across the enterprise by bringing data together in real time
- A unifying architecture for business rules and policies by providing a rule-driven architecture for the flow of information
- Architectural scalability for the enterprise by breaking performance dependencies between systems
- Architectural availability for the enterprise by breaking availability dependencies between systems
- New kinds of applications by exploiting real-time data and application integration

roles of the ZLE data store

convergent and integrated view of the enterprise

Typically, data in a company's operational systems is application specific or centered on a particular clerical function. Applications are created to address a specific business problem, product, or channel. As shown in figure 2, the same customer is assigned different identifiers in each of a company's operational systems.

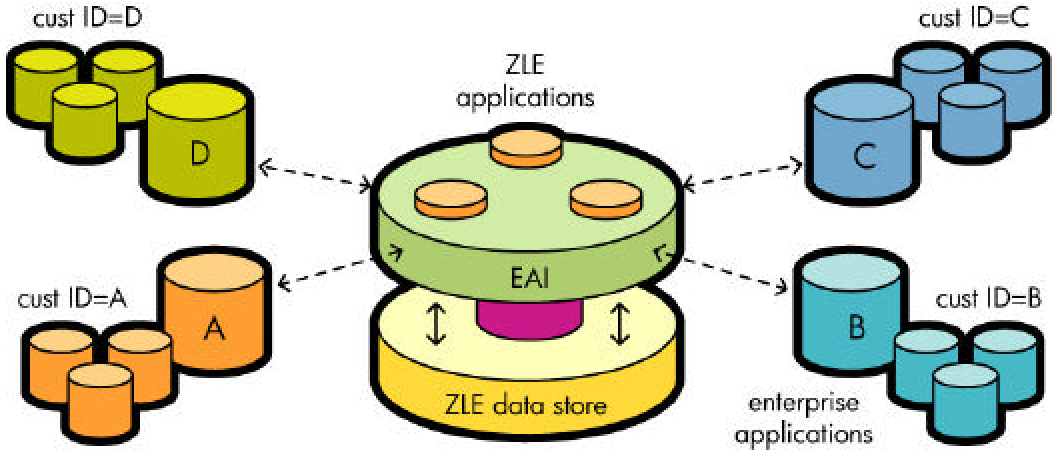

figure 2. application-centered enterprise data architecture

In a ZLE environment, on the other hand, the ZLE data store maintains a single, 360-degree view of the customer. Customer information is cleansed, stored, and enriched for distribution to all applications in the enterprise. Likewise, many additional subjects that are maintained in the ZLE architecture are consolidated into a single location. Thus, the ZLE data store represents the single view of the enterprise data set that serves as the basis for next-generation ZLE applications.

EAI message store

EAI is essential to a ZLE solution. With EAI, if a change is made to the data of one application area, then that change is propagated to other systems and databases throughout the enterprise or even outside the organization, for example, in business-to-business applications.

In figure 3, the ZLE data store acts as a message store. Messages are recorded in the ZLE data store as they are processed by the EAI hub. The EAI hub guarantees that revisions are propagated (mapped, routed, and applied), even if one or more of the systems cannot process the change when the EAI hub receives notification. For example, if the server that application B's database runs on is not available to process the change, then a record of that change is maintained in the ZLE data store. The EAI hub continues to attempt to process the change on application B until delivery or processing is confirmed. Some EAI transactions are very complex or require multiple steps to complete. In this case, the ZLE data store acts as a message store to track the "workflow" of the transaction until all execution steps are completed.

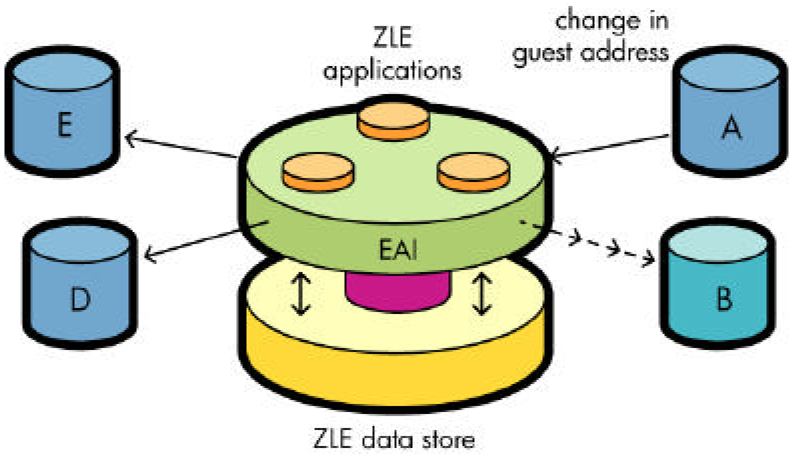

figure 3. EAI message store and memory

EAI memory and reduced reach-back dependency

Unlike traditional EAI, in which a change transaction or message is eliminated once guaranteed delivery occurs, the zero latency EAI has a memory. The transaction or state change continues to reside in the ZLE data store, reducing the need for EAI reach back. For example, in the classical EAI approach, if a new e-store ZLE application requires enterprise data, then it must either get and maintain another copy of the enterprise data or reach back transactionally to existing applications via the EAI hub to get the data. This poses several significant problems:

- The path length and response time of the transaction may increase prohibitively—e-store users will not wait very long before exiting to another website.
- The applications and databases that are being reached back into may not be able to handle the increased workload—they may be running at full capacity.
- The network bandwidth between systems may not support the additional workload.
- The applications that are being reached back into may not be online; thus the e-store transactions cannot be completed.

With zero latency data maintained in the ZLE data store, the e-store and any other real time–dependent applications can simply use that information when processing transactions.

detailed event and interaction knowledge store

In a ZLE environment, information retained in the ZLE data store must be very detailed. For example:

- For a telecommunications customer, the ZLE architecture needs a call detail record of every call made and possibly even uncompleted calls attempted.
- For a banking customer, the ZLE architecture needs a record of every loan payment, ATM withdrawal or deposit, credit card purchase, check, and more.
- For a retailer (brick and mortar or e-store), the ZLE architecture needs a record of every purchase line item, payment type used, and so on.

- For an airline or reservation system, the ZLE architecture may need a record for every flight operation, agent, or user activity, such as availability requests and segment changes.

Also, with the new customer touch points (Web, kiosk, ATM, and wireless), information in the ZLE data store must reflect interactions as well as transactions. Mouse clicks, wireless keypad touches, or voice-activated interactions offer valuable information on customer interests and preferences. Knowing what offers were made during each interaction and the respective responses gives a company invaluable insight that can be used to promote customer satisfaction and increase sales revenue.

state and lookup knowledge store

In a ZLE environment, two other types of information must be maintained by the ZLE data store: state and lookup data. *State data* conveys the current status of an entity. For example, it may describe the services to which a customer subscribes or current account balances. *Lookup data* conveys the approved or authorized members of a data set as well as their key attributes. Examples of lookup data include customers, current sales locations, region codes, and product and SKU numbers. While state and lookup data typically constitute a very small percentage of the total information in the ZLE data store, they are highly complex.

historical information store

The ZLE data store also contains a historical perspective or memory of information. This history is essential to evaluating prior customer behavior (transactions and interactions) and enterprise events and their impact on the organization and its customers. Only through an analysis of previous events, for example, by applying business rules, can the appropriate response be determined to an event or interaction that is occurring now. Historical information is also valuable in supporting instantaneous, enterprisewide information needs, such as handling customer service inquiries, settling billing disputes, or fulfilling legal or regulatory information requirements.

How much history is required? The time frames vary from business to business, but in general they are significantly less than the data warehouse requirements that typically store many years' worth of data. Essentially, it comes down to how much history is needed by the business rules engine to determine an appropriate response to an event or interaction. In some cases, that may be as little as 60 days. In other cases, in which customer events and interactions occur less frequently, several years' worth of data may be required. Another consideration may be how much information must be made accessible online for customer service and support. Regardless of the extent of historical data required, the ZLE data store must be designed so that the historical period can be extended easily without a redesign of the business or operational processes.

real-time data warehouse

The ZLE data store contains current and integrated information that cannot be found in any other data store in the enterprise, including the corporate data warehouse. As a result, it serves the real-time business intelligence needs of the enterprise. It also supports direct queries via industry-standard business intelligence tools. The ZLE data store can be used to extract information for batch processing, data marts, or even OLTP systems that need clean, current, and integrated data.

infrastructure demands of the ZLE data store

maintenance of current and volatile information

The ZLE data store requirement to maintain information that is as current as possible makes robust (in terms of both fault protection and bandwidth) links to operational systems essential. Indeed, the contents of the ZLE data store are constantly changing with the addition of event detail data and the maintenance of state and lookup data.

This amounts to a very complex set of I/O performance demands on the ZLE data store. It must be able to support, concurrently, the addition of hundreds, thousands, or even tens of thousands of events per second. Furthermore, these are not load-appends, as may be the case in a data warehouse environment. Because the information must be retrieved in a "transactional" fashion, random insertions are frequently used to achieve the appropriate clustering and subsecond access. Thus, issues such as block split management, index maintenance, and data space management emerge.

The ZLE data store must also support

- Massive random reads for zero latency transaction inquiries
- Constant in-flight updating and deleting of information in a random fashion to reflect current state or lookup changes
- Bulk deletions to remove aged data

All combined, the ZLE data store I/O workload is one of the most demanding asked of any database management system (DBMS).

massive concurrent user support

While a large data warehouse typically supports hundreds of users, the ZLE data store may be subject to demand by thousands or even hundreds of thousands of users. High-volume users include call center agents, Web guests, mobile phone access agents, event and interaction capture agents, EAI agents, and operational analysis or clerical users. Additionally, there may be some complex query users, but their workload is of lesser priority than that of transactional users.

optimized for subsecond interactions and transactions

The vast majority of ZLE applications are OLTP applications that require subsecond transactional access to the ZLE data store. This demand for subsecond access—combined with the need to support high-volume, concurrent users—requires

- Database schemas and a physical DBMS layout that are designed for extremely fast and random access; typically, a third normal form design is used to retain valuable entity relationships
- Transaction processing monitor control and access, that is, application services; these services are essential to support
  - Load balancing of applications across server resources

- Parallel execution control of application services (service creation and deletion)
  - Application monitoring and restarting in the event of failure
  - Transaction management
  - Reliable queuing
  - Connection management (condensing the number of DBMS connections)
  - Code sharing between multiple users of the same application, radically reducing the amount of storage required to support each signed-on user
- Transactional data integrity protection—two-phase commit protection across homogeneous as well as heterogeneous hosts
- Transactional data concurrency protection to maximize simultaneous access to the ZLE data store, while providing the appropriate level of data consistency required by applications; features such as “read-thru-lock” (dirty reads) provide fast and unencumbered access to information

tolerance of complex queries

While the primary design point of the ZLE data store is for transactional, subsecond use, the ZLE data store must also support concurrent access by complex query users on either an ad hoc or a predetermined, scheduled basis. Examples of these types of users are one-time business intelligence queries, data mart extracts for data mining or business intelligence, operational analysis reports, business rule deployment, and view creation. Furthermore, the ZLE data store must handle these complex queries without affecting the stringent service-level requirements of response-critical ZLE components.

In this mixed-workload environment, the highest priority activities are given first access to server resources, such as processor cycles, disk I/O, and communication bandwidth. Such an environment is complex because different vendors normally supply the system software, transaction processing monitors, applications, and DBMSs. Though each may priority-enable its specific piece of the solution, control of priority execution is lost when these components must interact with others. Additionally, mixed workload is complicated even more when clustered servers or massively parallel processing (MPP) servers are employed in a ZLE solution. Cross-cluster or cross-processor priority is difficult to control and manage.

around-the-clock availability

Customers expect to be able to conduct business when (day or night), where (from home or a hotel on the other side of the world), and in the form (call center access, Internet, or wireless) they choose. To compete effectively, companies must provide continuous processing.

When a company embraces a ZLE model for doing business, the ZLE architecture, with its ZLE data store, becomes business critical. If the ZLE data store is rendered inaccessible, then processing, as well as all of the applications linked to it, is terminated. As a result:

- EAI processing cannot be completed—synchronization stops.
- Event and interaction capture stop—these events may never be recaptured because of the volume at which they occur.
- ZLE applications that previously accessed the ZLE data store cannot be executed—these are applications that touch the customer, reduce fraud or risk, and generate revenue.

ZLE raises the bar on around-the-clock availability beyond just hardware fault tolerance. In this environment, the ZLE data store provides information even when operational systems are brought offline for batch reporting or database maintenance. It must remain online (read and write) during database maintenance operations (like reorganization and backup) and system configuration. According to some companies, it must even survive disasters, such as fires, earthquakes, and flooding.

filling the information gap

operational and transactional database gaps

Typically, operational database environments consist of stovepiped applications and databases and heterogeneous platforms and middleware (DBMSs, transaction processing monitors, and so on). They cannot fulfill the ZLE requirements for information because

- Transactional systems are fragmented, resulting in lack of a single view of the enterprise.
- Transactional systems are highly optimized for the most critical business transactions, typically deploying special data encoding and compression techniques.
- Data architecture is often complex and makes the use of user-friendly ad hoc query tools unacceptably difficult.
- Transactions have critical response times that would be adversely affected by unconstrained query use or additional EAI workloads.
- Traditional transactional systems can rarely tolerate the amount of historical data required in a ZLE environment.

data warehouse gaps

The ZLE framework differs significantly from a traditional data warehousing (EDW or data mart) architecture. Typically, an EDW or data mart architecture possesses the following characteristics:

- Relaxed extraction, transformation, and load (ETL)—usually loaded or refreshed nightly or weekly
- New data inserted or load appended—rarely, if ever, updated
- Support of strategic decision makers—not tactical or operational decisions
- Relatively few users performing large-scale, complex query analysis—exceptionally fast parallel queries required
- Response time measured in minutes as opposed to transactions per second (TPS)
- Accessed via business intelligence, OLAP, and data mining tools as opposed to applications managed by transaction processing monitors
- Database designed to enhance intraquery parallelism and provide simplified end-user access—frequently a denormalized structure such as a star or snowflake schema
- A consumer of enterprise data as opposed to an integrator of information—no EAI performed; it receives data from legacy operational systems and may feed data marts, but there is no true EAI as a primary mission

- Not typically 24 x 7—downtime required for ETL processing, database maintenance, and creation of aggregates and stable views

As discussed previously, the ZLE model requires a data environment (typically normalized data for fast transactional access and updating) and transaction and query execution environment that are entirely different from the EDW or data mart. In a ZLE environment, data (event detail, state, and lookup) should flow into and out of the ZLE data store via EAI with no or low latency. Data is constantly updated to remain current with the operational systems. The ZLE data store is used for making real-time decisions that affect customer service and customer relationship management (CRM), e-business (or m-business) activities, and operations. Hundreds of thousands of these decisions, which typically require subsecond response, may be executed per second. Changes resulting from ZLE applications must be propagated to operational systems to keep the enterprise synchronized.

information gaps

Another way of looking at the problem with existing data frameworks is from an information gap perspective, that is, data needed to fulfill ZLE requirements. In a paper titled *Information for the Real-Time Enterprise* (Buckinghamshire, U.K.: Bloor Research, 1999), Bloor Research developed information gap truth tables that looked at the various information needs of an enterprise. Their report considered a matrix based on all combinations of query dynamics (see table) and analyzed whether or not the information could be provided by existing operational systems or by the data warehouse.

truth table

transaction / query dynamics	*transaction / query complexity*	*transaction / query volume*	*data currency*	*data level*	*source system*
predefined, ad hoc	simple, complex	low, high	real-time, historical	transaction, aggregated	integrated, fragmented

This information gap analysis determined the total combination of query dynamics that were satisfied by a data warehouse or existing transactional data stores:

- 9 of 32 combinations of real-time information
- 9 of 32 combinations of transaction data
- 10 of 32 combinations of fragmented systems
- 11 of 32 combinations of ad hoc queries
- 11 of 32 combinations of complex queries
- 12 of 32 combinations of high-volume queries

As can be seen, more than half of the possible combinations of transaction and query parameters were not satisfied by data in the transaction processing systems or against a traditional data warehouse. The worst-case scenario was

- Transaction / query volume: high
- Data currency: real-time

- Data level: transaction
- Source system: fragmented

This listing profiles the zero latency information needs of an enterprise. The role of the ZLE data store, however, is to bridge the information gap by providing an environment tuned to information delivery, containing data at the event detail level, coordinated across all relevant source systems, and maintained in a current state. The result: an integrated and up-to-date view of the entire business.

positioning the ZLE data store in the enterprise architecture

One way that the ZLE data store can fit into your enterprise architecture is shown in figure 4. In this model, the ZLE data store is the "leading edge" to the corporate EDW. Using EAI technology, it can feed data to the EDW in trickle mode or in prescribed batch feeds.

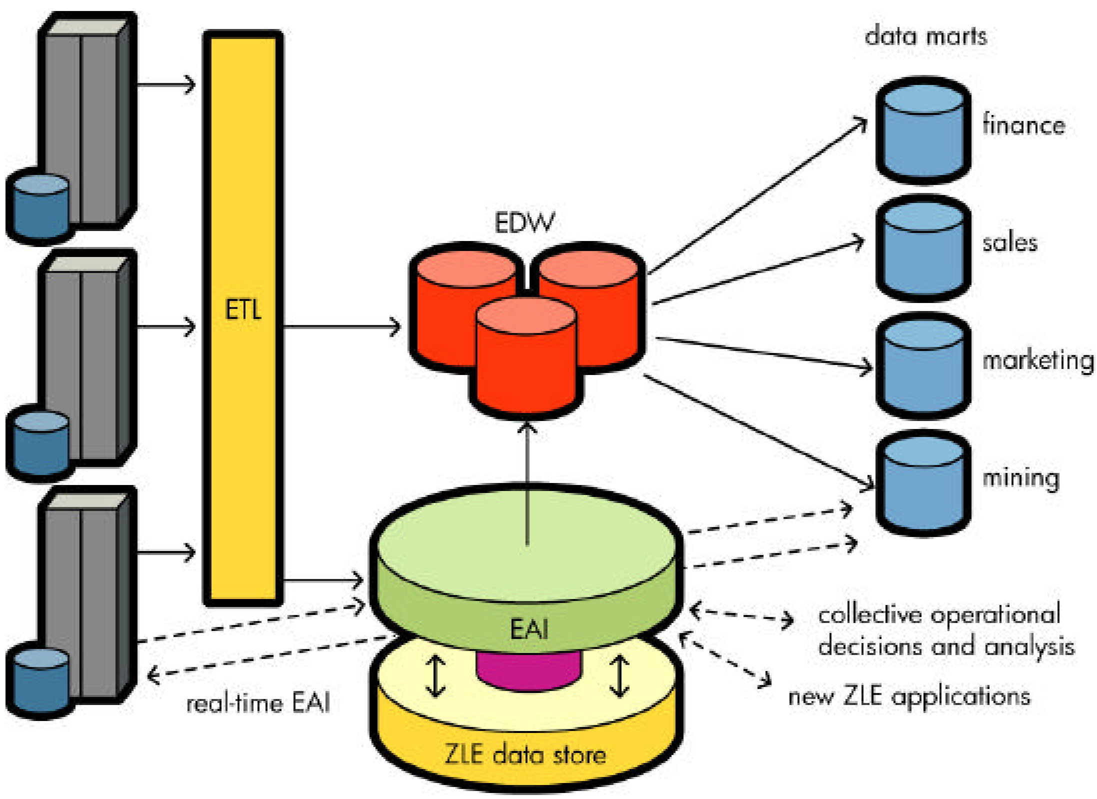

figure 4. deploying the ZLE data store—model 1

An alternative model (see figure 5) can be deployed in which the ZLE architecture fulfills the expanded role of the EDW. In this model, the ZLE data store supports the direct creation of data marts rather than creating them from the EDW. Business intelligence or data mining analysis (for example, for fraud, churn, or propensity-to-buy predictions) can be performed on mining extracts of the ZLE data store to develop new business rules and policies. These rules are employed immediately when they are placed back into the ZLE data store.

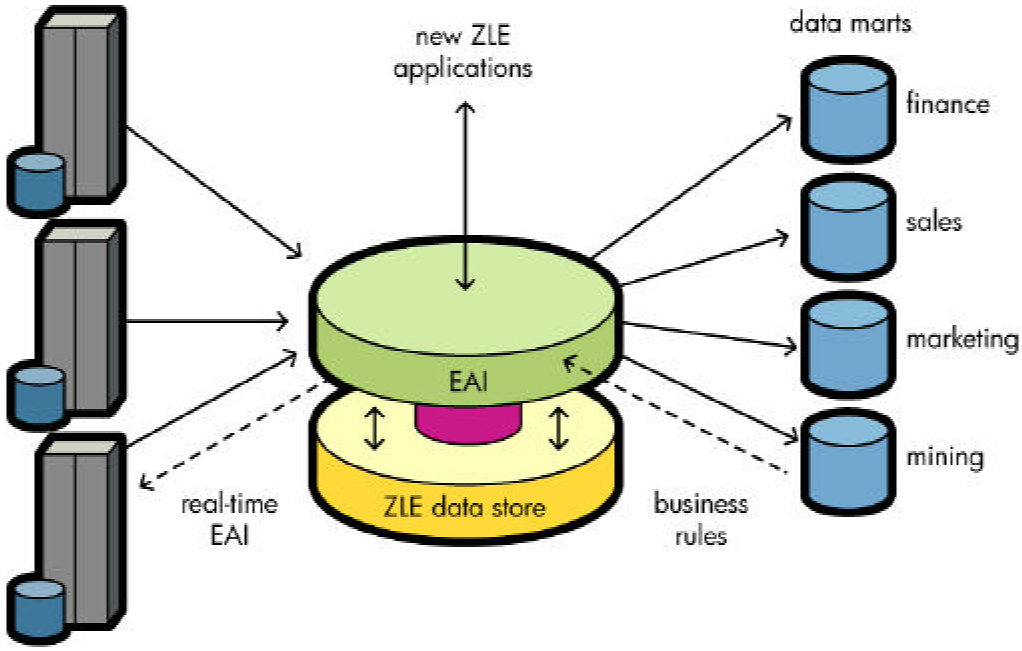

figure 5. deploying the ZLE data store—model 2

ZLE data store technology challenges

application and database scalability and parallel performance

In a zero latency enterprise, applications and databases do not have to be redesigned or reengineered to support additional applications or users or a growing database. The prevalence of mergers and acquisitions and the move from transactions to interactions (clickstreams or wireless key interactions) make database scalability and parallel SQL capabilities essential. Also, scalability must extend to all components, including hardware and middleware. If all components do not scale equally, then the solution cannot grow to meet the challenges of a ZLE environment. This is the fundamental reason why "shared-nothing" or MPP designs are the preferred architecture for the ZLE data store rather than shared (SMP and shared-memory) architectures.

business-critical reliability and availability

In a zero latency enterprise, applications and databases must be available when the business or customer demands them. This can imply support for 24 x 7 applications, or it can simply mean that during the hours of service no outage is acceptable. It can also imply business continuity (disaster recovery) protection of the ZLE data store. All elements of availability must be considered, including hardware, system software, middleware, and even operations.

manageability of very large tables and databases

Very large database (VLDB) manageability is the ability to monitor, configure, and initiate maintenance command and control over very large (multi-terabyte) database objects:

- *Online* (no interruption of service, including full read, write, and update during operation)
- *Globally* (no matter where the object is running—in one or multiple distributed servers)
- *With fewer resources* (automated operations and single system image)

With VLDB objects, database operations must be performed both in parallel and at a granular level (down to the individual partition) to reduce the time required to complete the operation. For example, system overhead costs (processor and I/O) and time to complete are both dramatically reduced by reorganizing only a two-gigabyte partition of a table as opposed to the entire terabyte table.

complex and high-volume mixed-workload management capabilities

Very large ZLE environments must support, concurrently, many different users and applications, which possess many different I/O profiles and priority requirements. This mixed-workload, parallel environment must grant priority access to system resources consistent with enterprise business goals. For example, a mixed-workload environment might include *simultaneous* support for

- Large-scale and real-time ETL requirements
- Large-scale OLTP or operational access detection, real-time reporting, and real-time data mining
- Batch subscribers, such as data mart extracts and EAI data feeds
- Complex queries, such as OLAP and complex reporting
- VLDB maintenance, such as reorganizations, backups, and moving data from one disk to another to alleviate hot spots
- Loading and updating events and state changes into the ZLE data store continuously without disrupting service levels
- Providing access from the ZLE data store to a very large archive of detailed data
- Robust integration of ZLE middleware components

hp ZLE capabilities

HP has uniquely sourced and integrated technology for creating robust ZLE solutions, including the hardware platform, DBMSs, transaction processing monitors, EAI middleware (transformation, workflow, adapters, and routers), and rules engines. Many of these technologies are so interrelated that the technology infrastructure has to be considered as a whole rather than in parts, but it is important to understand the roles of the most important components.

The foundation of HP ZLE solutions is the reliable, scalable NonStop platform, which is uniquely capable of meeting the technical challenges of delivering on the ZLE data store promise.

proven linear scalability and parallel performance

The NonStop server offers virtually unlimited processor scalability—from 2 to over 4,000 processors. The I/O subsystem scales dramatically as well in terms of I/O bandwidth, disk drives and controllers, and network I/O components. This scalability is enabled by HP's unique shared-nothing architecture, operating system, and systems services, for example, database, transaction processing monitors, communications, and database inserters. Each processing element is a processor with its own associated memory, I/O, and copy of the operating system. Processors execute independently and in parallel. NonStop software is designed to leverage this multiprocessor architecture to provide massively parallel execution. As a result, each processor adds a full processor's worth of power, unlike symmetric multiprocessing (SMP) configurations. This is the key to the linear expandability and scalability of the ZLE data store based on the NonStop server.

Like all system software for the NonStop server, the NonStop SQL database is a shared-nothing DBMS. Its architecture allows it to benefit nearly 100 percent from the raw resources of any hardware component added to the loosely coupled server. No extra processing is required to take new hardware into account, since each node is fully independent of the others. NonStop SQL software does not rely on a distributed lock manager (DLM). Each data access manager (DAM) has full and exclusive control (locking, cache management, and logging) of a subset of the database partitions. As a result, the NonStop SQL database can provide unprecedented linear scalability and parallel performance as processors and disks are added to the ZLE data store server.

Massively parallel application services are transparently enabled through HP application services software, including NonStop Tuxedo, NonStop CORBA, and NonStop Server for Java™. Applications developed for other platforms can automatically take advantage of the MPP resources of the NonStop server without any platform-specific coding. By leveraging the MPP capabilities of the NonStop platform, companies can achieve unprecedented levels of transactional performance and scalability.

business-critical reliability and availability

NonStop servers enjoy the number 1 ranking for industry-recognized leadership in availability. This is made possible through the unique combination of hardware and software fault-tolerant features built into the NonStop platform and system software. Patented process-pair technology and middleware permit objects and transactions developed for other platforms to inherit NonStop characteristics without specialized application code.

Essential components of the NonStop SQL database are developed using HP's patented process-pair technology to enable continuous availability. If one processor fails, within a fraction of a second a DAM running on another processor automatically takes over all of its data access tasks. As a result, the end user experiences no interruption of service. The NonStop SQL database uses the NonStop platform's efficient and reliable message system to communicate between subtasks, ensuring optimal performance and reliability. In addition, it is integrated with system management tools to minimize the need for operations staff while maximizing performance and availability.

NonStop Remote Database Facility (RDF) software provides business continuity and replication services on behalf of the NonStop SQL database. NonStop RDF software asynchronously pulls completed transaction afterimages from the NonStop Transaction Management Facility (TMF) transaction log and provides guaranteed delivery and replication of key ZLE data store database tables to one or more remote sites. It is designed with the scalable high performance and superior reliability demanded by the ZLE data store. HP's major financial and telecommunications customers use NonStop RDF software extensively to enable 24 x 7 availability of business-critical applications and data.

manageability of very large databases

In addition to fault tolerance of the platform and software infrastructure, the NonStop server leads the industry in online manageability of the platform, network, or communications subsystem, database, and applications services environment. Examples of online DBMS operations include

- Back up partition, table, disk
- Reorganize disk
- Partition split, merge, move, move boundary, drop
- Create/drop/rename table, index table, view, constraint
- Reconfigure disk cache
- Roll forward partition, table, disk
- Selective logical data corruption recovery by table(s), transaction ID, column(s), transaction group, times, I/O operation
- Add column to populated table
- Add new disk or processor

The NonStop SQL database supports the management of massive, multi-terabyte data stores with highly parallel and granular management operations performed at the table or index partition level. These partition-level operations are executed in parallel, dramatically reducing the time required to complete the operation.

mixed-workload support

NonStop servers host many of the world's largest data warehouses, 90 percent of the world's stock exchanges, and the largest message-switching applications. Companies use NonStop servers to build massive, hybrid EDWs and ODSs that operate in a mixed-workload environment. These systems operate with predictable service levels; specifically, they sustain transactional inserts, batch extracts, OLTP, and massively parallel queries against the same tables simultaneously without degrading service levels. The NonStop SQL database stands alone in the DBMS industry for its ability to handle all types of prioritized workloads executing concurrently on the same server, despite very different I/O profiles.

prioritized I/O support

NonStop software is tightly integrated with the NonStop platform, enabling tight control over the mixed-workload environment. All components can plug and play and work together cohesively. For example, the NonStop SQL database overcomes issues associated with mixed workloads executing concurrently on the same server. Priorities can be assigned to individual workloads (queries, transactions, and management commands) based on the urgency of the information or the expected execution cost in terms of I/O and processor cycles. The priority of an individual workload dictates the processor cycles and I/O allocated to it relative to other workloads executing at the same time on the same server. This priority is carried with each resource request made on behalf of the query or application, even when requests are passed to other processors in the same loosely coupled MPP server.

Disk I/O requests are queued to the respective DAMs in priority sequence. When they reach the top of the DAM queue, an I/O request is processed only if it is the highest priority request executing in the processor. Thus, response rates for the most demanding transactional events are guaranteed, with concurrent support for large batch loads or trickle updates, online maintenance operations, data mining, or other complex analysis.

This priority mechanism, which is unique to HP, offers a level of mixed-workload support that no other DBMS vendor can match.

query control and management

The ZLE data store can support many lower-priority business intelligence queries that could be asked alternatively of a data warehouse or a data mart application. But its primary mission is to supply ZLE-enabled applications with real-time, integrated information. The HP ZLE framework has a unique mixed-workload capability that permits these lower-priority business intelligence queries to run concurrently with the ZLE workload. But they are constrained to utilize only “white space” resources (processor cycles and disk I/O) that are not required to meet the demanding service-level requirements of ZLE-enabled applications.

The ZLE data store based on the NonStop SQL database is query accessible from virtually all business intelligence tools or applications that support an Open Database Connectivity (ODBC), ADO, OLE DB, or Java Database Connectivity (JDBC) interface. HP provides enhanced query control through priority management built into gateway products such as NonStop ODBC Server software. Each user or group is assigned a unique profile that indicates the priority scheme for executing SQL statements based on the optimizer’s determination of the expected cost (processor cycles, I/O, and message costs). Operational queries (queries with low estimated processor and I/O cost) can be assigned higher priorities automatically, while complex queries can be assigned lower priorities so that they execute as background tasks.

NonStop ODBC Server software supports a configurable preemptive control for each user that prohibits the execution of queries that have an estimated cost exceeding those controls. NonStop SQL software monitors queries being executed and can alter their execution, lowering their priority, or stop them if they perform excessive I/O.

HP is the only vendor to provide this degree of sophistication in query control and management in a mixed-workload environment.

robust integration of middleware components

HP ZLE solutions are designed to help companies by integrating the multiple components needed for EAI, database management for the ZLE data store, integration adapters, transaction processing monitors, and ZLE application creation and hosting. The tighter and more robust the integration and the fewer solution pieces involved, the lower the risk and cost of ownership. HP software for the ZLE environment includes

- NonStop SQL
- NonStop ODBC Server
- NonStop CORBA
- NonStop Server for Java

- NonStop Tuxedo
- NonStop TMF
- NonStop RDF
- NonStop Data Transformation Engine
- NonStop DataLoader/MP
- NonStop dba/m
- SMF

In addition, HP has established key ZLE partnerships with independent software vendors (ISVs) that can provide added value to a ZLE solution, such as EAI adapter technology, business rules engines, business intelligence, OLAP, data mining tools, and data mart extraction tools. These partnerships include extensive integration testing, Software Development Kits, attractive pricing packages, and services. As a single point of contact, HP Services enables access to any and all ZLE ISV partner applications.

ZLE demonstrations

In partnership with HP, many companies have embarked on implementing ZLE solutions. Because of the competitive advantages gained, many companies are reluctant to share their experiences and return on investment (ROI) results.

In response to requests to see ZLE in action, HP has established several of the world's largest demonstrations (111-terabyte telco and 50-terabyte e-CRM). These demonstrations showcase the capabilities of HP ZLE solutions, in which enterprise application integration and data integration are merged through a real-time ZLE data store.

telco ZLE demonstration

The telco demonstration shows the following components running simultaneously:

- 111-terabyte ZLE data store composed of over 100 billion transactional events
- ETL processing of over 1 billion transaction events per day and over 14,000 per second (random inserts)
- 40,000 call center customer service reps running over 3,000 TPS, with one-tenth of a second response time each
- 100-TPS EAI delivery of events to a credit data mart running Oracle® software
- Very large ad hoc OLAP queries executing against the ZLE data store
- Real-time publication and subscription for event alerting
- Batch subscription for extracts to data marts

ZLE e-CRM demonstration

HP recently constructed a proof of concept (POC) demonstrating a ZLE solution for supporting real-time e-CRM. The solution provides a complete customer view, integrating customer information from all operational systems throughout the enterprise in a central ZLE data store. It features the following:

- 50-terabyte ZLE data store composed of point-of-sale and e-store clickstream and purchase detail activities

- Complex mixed-workload transaction mix composed of simultaneous event capture, EAI, OLTP, relational online analytical processing (ROLAP) queries, and data mart extracts
- Multiple data channels (touch points) with events generated at variable rates by driver systems
- Data cleansing and de-duplication functions
- Incorporation of ROLAP and real-time event notification tools
- Call center software
- Campaign management tool
- Integrated data mining tools
- Data staging tool
- Customer Manager, developed by HP, including customer demographics with real-time data enrichment
- Interaction Manager, developed by HP, which is a ZLE service application that utilizes business rules to respond to requests from cross-enterprise touch points, allowing offers to be made to customers based on customer profiles and events. Offers made are stored as events in the ZLE data store, ensuring that all touch points have knowledge of customer interactions from other touch points

As a part of this POC effort, HP developed a ZLE framework for retail solutions. HP is prepared to tailor and integrate the framework into a comprehensive e-CRM solution to meet the complex needs of any company.

complete life-cycle services

HP Services and HP consultant/systems integrator (C/SI) partners offer complete life-cycle services to help organizations implement ZLE solutions rapidly with minimal risk and operate them continuously.

- ZLE Planning Services
- ZLE Design Services
- ZLE Implementation and Deployment Services
- ZLE Management Services
- ZLE Business-Critical Support Services
- ZLE Business Continuity Services

conclusion

At the heart of any ZLE solution is the ZLE data store. The ZLE data store bridges the information gap between operational systems and the data warehouse. It provides a data store tuned to information delivery, containing data at the event detail level, coordinated across all relevant source systems, and maintained in a current state. Thus, the ZLE data store provides the integrated and up-to-date view of the business that is essential for a zero latency enterprise.

In a zero latency enterprise, the ZLE data store must be tightly coupled with EAI, as it is in the HP ZLE framework. The ZLE data store fulfills many complex roles; consequently, it presents an extremely demanding set of infrastructure requirements and technology challenges. As the basis of the ZLE framework, the NonStop server and system software provide the reliability, availability, and scalability to deliver on the ZLE data store promise.

HP has the technology, ISV and C/SI partnerships, and services necessary to support the world's largest and most complex ZLE environments. HP customers and demonstrations prove that a robust ZLE solution can be built and that real business benefits can be achieved. As Roy Schulte, ZLE business practice manager for Gartner, notes, "By integrating applications and the high-availability hardware, zero latency enterprise can cut up to 40 percent from the time, money, and resources companies would need to do the installation and integration themselves. Every project comes with a lot of risk. Compaq's [now HP's] ZLE can reduce the risk of failure."

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

Printed in the U.S.A. 02-0388

invent