a white paper
from hp

Real-Time ODS Manager

Enabling the Real-Time Enterprise

Abstract: In an information-based environment, organizations that eliminate "information float"—the costly lag between the time data is created and when it gets used—have a natural advantage. Accelerating the collection and use of decision-making data is increasingly a critical requirement.

Recognizing the advantages of eliminating latency, Hewlett-Packard has developed a unique capability that combines a high-performance operational data store (ODS) with enterprise application integration (EAI) technology. The result is a real-time enterprise solution that can handle high-volume mixed workloads in an Oracle database environment. A mixed-workload environment enables the continuous insertion of records and the simultaneous processing of queries and transactions in support of real-time decision-making and business processes.

This ODS solution runs on *AlphaServer*™ systems in the *Tru64*™ UNIX® and HP-UX operating environment. It uses the Oracle*9i* Real Application Clusters (RAC) database, advanced HP clustering technologies, newly developed Real-Time ODS software, and innovative ODS design and consulting services. These capabilities enable solution providers and customers to build solutions that address the evolving needs of the real-time enterprise today.

This white paper explores the concept of the operational data store and outlines the history of HP's development of Real-Time ODS capabilities as the foundation of Real-Time Solutions for Oracle databases. It concludes with guidelines for assessing whether an organization could benefit from an ODS and, if so, how HP and its partners can assist in implementation.

Table of Contents

1. Executive Overview

> "The agency does not have an effective system to analyze important pieces of information that come from different sources at different times but should be linked."
>
> FBI Director Robert S. Mueller III, *The Boston Globe*, May 30, 2002

Director Mueller's observation quoted above highlights one of the greatest problems with information – there is too much of it in too many places. From government agencies to telecommunications firms to banks and retail operations, managers and analysts are handcuffed and blindfolded by their inability to find the right answers at the right time. And in this always-on world, the right time is all the time, any time.

At HP, we invent IT solutions to master the information environment. Our innovations, developed through close industry partnerships, are ready for you to incorporate into an *always-on enterprise infrastructure.* This is an infrastructure that you can design effectively, manage seamlessly, optimize dynamically, evolve intelligently, and leverage operationally to meet your business goals.

This white paper focuses on the concept of leveraging an infrastructure with an operational data store (ODS). An ODS puts detailed, current information from disparate sources at the fingertips of the organization. It enables a database to sustain a high data-insertion rate while simultaneously enabling online transaction processing (OLTP) and decision support system (DSS) queries on that data. Contrast this with a data warehouse where the information contained therein is a slice of time. Valuable though it is for analysis, the data from this snapshot presents a static picture of the past, not a live image of the present.

An ODS is in a constant state of update providing the most current data for business operations. Now the left hand can know what the right hand is doing – and, even better – help it do its job. Access to real-time information from an ODS supports and enhances business intelligence, risk management, customer relationship management, forecasting, marketing campaigns, and any number of organization-specific information-driven activities. An ODS

- Enables real-time decisions based on up-to-the-second information
- Provides a real-time, integrated view of enterprise data
- Helps you close the loop on business processes
- Allows you to manage growing volumes of information in real time

An ODS clearly has great business value. This white paper explores that value as well as the history and technical framework of the HP Real-Time ODS Manager. Furthermore, the paper will discuss whether an ODS is right for your business or organization and, if so, what steps you can take to put one online.

2. Business value of an operational data store

There is a lot to be said for getting the answers to “what-if” questions in real time. Many businesses and organizations face challenges in today’s hyper-informative environment that require action in real time in order to maintain a competitive advantage.

Today, businesses and organizations are limited in their ability to act on real-time information. One limitation is the inability to create an integrated view of information—to have a 360-degree perspective of the customer or the enterprise. Another limitation is the inconsistent application of business policies. You might treat customers differently depending on whether they phone your call center, walk into your retail store, or interact with you on the Web. Other limitations include a lack of consistent information quality and the inability to integrate information from disparate databases throughout the enterprise or extended environment.

Where real-time solutions are required

In general, real-time solutions are deployed to recognize patterns using business intelligence. This enables a business or organization to act appropriately and quickly whenever critical patterns occur. Some specific examples follow.

Banks need to immediately recognize trading situations that involve excessive risk. With an integrated view, they can manage the risk in real time.

Telecommunication service providers need to prevent fraud in real time. When a calling card is stolen, it might be used to call all over the world on the service provider’s dime. It is critically important to detect an abnormal pattern, recognize it as fraud, and stop the fraud in real time.

And probably the most important need, relevant to all businesses, is the ability to acquire and retain customers. That means knowing every customer regardless of how they interact, and behaving intelligently every time they are in contact.

Manage risk. Combat fraud. Keep customers happy. An ODS can help.

Close the loop on your business processes

HP Real-Time ODS capabilities provide an always-on, real-time view of enterprise information. It enables the insertion of new data while simultaneously supporting queries and OLTP against a single clustered database in real time. It allows an organization to close the loop on its business processes with a fast and assured response time against a high volume of transactions.

3. Development history of Real-Time ODS capabilities

Conventional wisdom holds that mixing high-performance OLTP and DSS queries with a high volume of insertions on the same database is to be avoided at all costs. Yet, that is the defining feature of an operational data store. In the quest to develop a real-time ODS, HP engineers questioned conventional wisdom. They conducted an investigation to determine how well the Oracle shared-disk database engine could support a mixed workload. They wanted to create an environment that would enable

a large number of users to query data that is refreshed simultaneously and continuously with a high volume of insertions.

What follows is a short history of the research, testing, and subsequent development of new techniques to support a high-volume mixed workload predictably. The new techniques -- the subject of a recent patent application -- have produced a three-fold improvement in mixed-workload performance.

Feasibility study

The introduction of partitioned tables in Oracle Version 8 was the starting point. Engineers in our advanced solutions engineering center realized that it might be feasible to use the Oracle database as a mixed-workload storage element in a real-time architecture.

To this end, they conducted a feasibility study with a model database similar to one that might be used by a typical Web-based retailer. They used this prototypical database structure and an eight-processor SMP system to examine the effect of table partitioning on performance. The results encouraged them to invest more heavily in an engineering program to investigate what else might be done to develop a high-performance, mixed-workload environment. This program comprised three major phases over the past two years.

Phase 1 – exploring the domain

The object of this phase was to determine whether an Oracle shared-disk database could support enterprise-scale performance in a real-time environment at enterprise data volumes – thousands of inserts and hundreds of queries per second. Phase 1 went beyond the feasibility study to a larger hardware environment to find the maximum insert performance levels attainable while simultaneously executing a simple DSS workload. It provided the opportunity to study the behavior of the system under a mixed workload and to use system clustering to further partition the workload for better performance. At the same time, the engineers worked to uncover any unforeseen limitations that could affect the availability profile at very large database sizes.

Phase 1: configuration

The Phase 1 configuration included a two-node cluster of 32-processor SMP servers with 1.7 TB of Fibre Channel attached storage; the *Tru64* UNIX operating system; TruCluster Server software; and Oracle version 8.1.7 with first one and then two instances, both using Oracle Parallel Server. The database was initialized with 50 GB of raw data, and a total tablespace of 275 GB. As data was inserted during the test, it grew to approximately 1TB. The data model was a snowflake schema with multiple fact and dimension tables that could be partitioned on line. Summary tables were created to satisfy frequently requested summary data or frequently used table joins.

Phase 1: workload description

The inserts to the selection table of the database were created as 60-byte records generated by a series of drivers on the system under test. The records were then batched into 100-record groups to be applied to the database. The queries were defined and executed from an external client driver program performing SQL procedures. By allowing multiple streams of the same set of queries to touch different

parts of the database, undesired memory caching effects were avoided. The queries used in were set to select a customer randomly and then get the list of products that the customer purchased in the last three months including records added in the last few seconds.

Phase 1: test methodology

The number of stimulus client processes was increased and the rate of inserts was measured along with CPU activity, memory utilization, disk/controller, and process/thread activity. The test was performed with and without query activity to measure the impact of a DSS load on the volume of inserts that could be processed. After each run, several tuning parameters were adjusted to determine their impact and then the database was reset.

Phase 1: results

Sustained insert rates were measured at 24,000 inserts per second (ips) on a single OPS instance (one 32-processor node) with over 100 DSS queries per second (qps) executed simultaneously with a query response time of 0.045s or less.

28,000 ips were observed without any DSS load. The maximum CPU activity only reached 52% indicating that further exploration was justified to determine where greater throughput could be achieved.

Phase 1: conclusions

The testing validated the theory that it was possible to handle a very high volume, albeit relatively simple, mixed workload with a shared-disk RDBMS. To put the insert rates into perspective, 24,000 inserts per second might represent the total call volume of a large national telephone company or perhaps five times the point-of-sale volume of a large national retailer. This, along with the relatively modest CPU utilization, gave the engineering team confidence that there was adequate “head room” to employ other techniques or to handle the type of workload that might be encountered in a real customer project.

Phase 2 – extending the environment

Phase 2 focused on increasing the complexity of the workload to more accurately simulate the real-world environment and to overcome unwanted interactions between workload elements observed during Phase-1 testing.

The following six objectives were set for Phase 2:

- With a more realistic workload, demonstrate that interactions between the insert and query domains can be segregated to the point where a level of service can be sustained for the incoming data records
- Demonstrate a smoother insertion rate
- Determine maximum OLTP rates for specified system configurations while simultaneously executing a fixed DSS workload with the new mixed-workload enhancements employed
- Determine the effect of varying DSS activity on transaction throughput
- Determine practical maximum values for DSS queries
- Demonstrate fail-over capability and understand any benefits and limitations stemming from the choice of Oracle in a clustered environment

Phase 2: test platform

The phase 2 configuration included a two-node cluster with each node comprising four 500 MHz CPUs and 16 GB memory (4 GB/CPU). The operating system, Oracle versions, and the database itself were unchanged from Phase 1.

Phase 2: workload description

The workload was adjusted to increase the complexity and type of queries. In addition to the array inserts of detailed records to the database, the transaction that reviewed customer history was enhanced by implementing two types of queries. The first was an update that extracted customer information from multiple tables and then combined that information to write to multiple tables. The second was a "heavyweight" query that gathered a list of all of the products a customer purchased over time, up to and including purchases made during the last second. The engineers stipulated that 90% of the transactions meet stringent response times.

Finally, a query was added that selected information from three tables based on customer ID number. This query returned information about the customer as well as a description of the last product purchased. It was stipulated that this third type of transaction not have more than a 10% impact on the customer history queries and must run for at least a two-minute interval in order to measure the impact.

Phase 2: test methodology

As in phase 1, the team measured the insert rate over several test runs to show the effect of increasing the transaction rates and varying the composition of the mixed workload. They then added a variable amount of "think time" between queries to understand the effect of throttling that component. Limitations in the I/O configuration prevented the system from being tuned for optimal performance though the functional impact of the various workload components could be observed.

Phase 2: results

The tests showed that increasing the query load by a factor of ten reduced the insert rate by a factor of two. The insert rate reduction then leveled out above 5,000 ips with 200 qps running simultaneously as shown in figure 1.

This indicated that the insert rate could be maintained above an acceptable level, in this case 5,000 ips. This achieved the desired sustained-level-of-service goal. By varying each component of the mixed workload, it was observed that the major culprit in terms of impact to insert rate was, as expected, the increase in heavyweight queries.

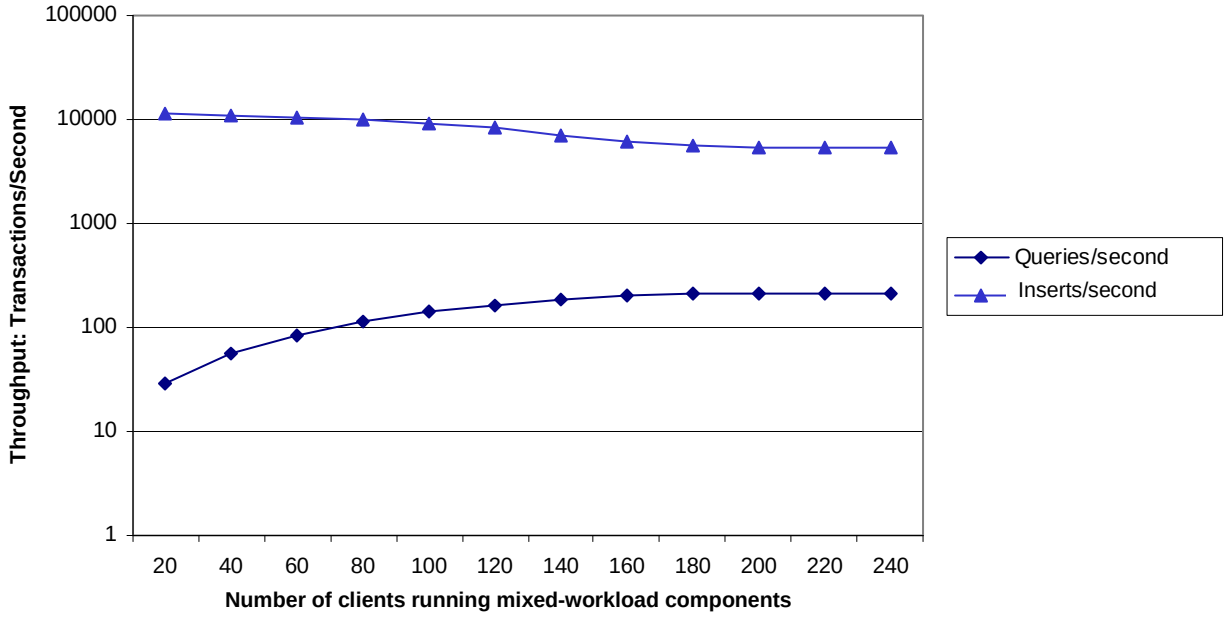

Figure 1 Throughput summary

Phase 2: observations and conclusions

Phase 2 clearly demonstrated that the team needed to design a resource manager to enable control over of the queries to reduce interactions in the database and operating system. In this scenario, the I/O was the bottleneck, not the CPUs. This could be fixed easily by reconfiguring the environment to use more appropriate storage devices and designing the database for minimal inter-instance traffic -- the amount of data being sent back and forth between Oracle instances. This phase resulted in the development of resource control software and the invention of an extension to the Oracle exchange method that is the subject of a recent patent application.

Phase 3 – proving the concept

Phase 3 featured the testing of the prototype Real-Time ODS Manager. It was aimed at acquiring the information necessary to size and characterize hardware configurations to meet specific project needs and, if possible, to derive a formula to describe the behavior on a given configuration. In addition, the engineering team wanted to compare database performance with and without the Real-Time ODS Manager.

Phase 3: test platforms

Two configurations were chosen based on assumed typical requirements with identical storage. All testing was done using *Tru64* Unix V5.1A, *Tru64* TruCluster V5.1A, and Oracle9*i* Real Application Clusters. It was further assumed that results could be extrapolated to version 5.0/Oracle8*i* from earlier test results, which showed a 15% improvement from 9*i* versus 8*i*. The chosen configurations were:

- 2-node cluster of 4 CPU machines
- 3-node cluster of 4 CPU machines

Phase 3: workload description

The test database was the same as that used in Phase 2 with small changes to reflect the Oracle9*i* environment. In addition, the team improved the physical layout and incorporated knowledge gleaned from the previous tests such as putting pairs of redo logs on separate devices and relying more on the disks for striping.

The team used the workload enhancement defined in Phase 2 plus some additional refinements as described below and used the same sub-partitioning and exchange features. They also used a similar formulation of mixed workload with minor changes and the addition of a third query class to closer approximate a real-world scenario. The queries ranged from simulated billing lookups and customer activity analysis to program usage profiles and daily sales activity reports. These different query types were characterized as short, medium, and long.

The short query simulated an eCRM query that returned all of a random customer's purchases for the last month. This query scanned the table currently being updated by the insertion program for the very latest data and ran without think time since it was simulating a software-driven event.

The medium query listed the products that a random customer purchased between two arbitrary dates and the time of day that the sale took place. This query simulated a customer-service lookup.

The long query simulated a long-term trending report that checked current buying patterns. It did a select of the top-20 best-selling products necessitating a scan of all of the sales of all of the products in the system for a given period.

Phase 3: test methodology

The number of inserts per second for each test point was again the salient data with the systems running at their maximum insert rate as the DSS load was varied. The team tried to determine the maximum DSS performance with no insert activity and recorded several data points in between. Because the systems were tested with the ODS Manager installed and the ODS Manager supports both the C++ CORBA client and the Java CORBA client, both interfaces were tested. The focus was on the higher performing C++ CORBA interface. Other data was collected as in the previous phases.

Phase 3: results

Using the ODS Manager in a mixed-workload real-time Oracle9*i* RAC environment, the team observed a three-fold improvement in terms of the number of inserts per second compared to an environment without the ODS Manager. There was an approximately three-fold increase in the number of queries per second achieved.

For example, as shown in figure 2, using the ODS Manager the workload was able to execute over 500 qps while inserting 15,000 new records into the ODS per second. Another data point showed over 600 qps while inserting over 9,000 new records into the ODS per second.

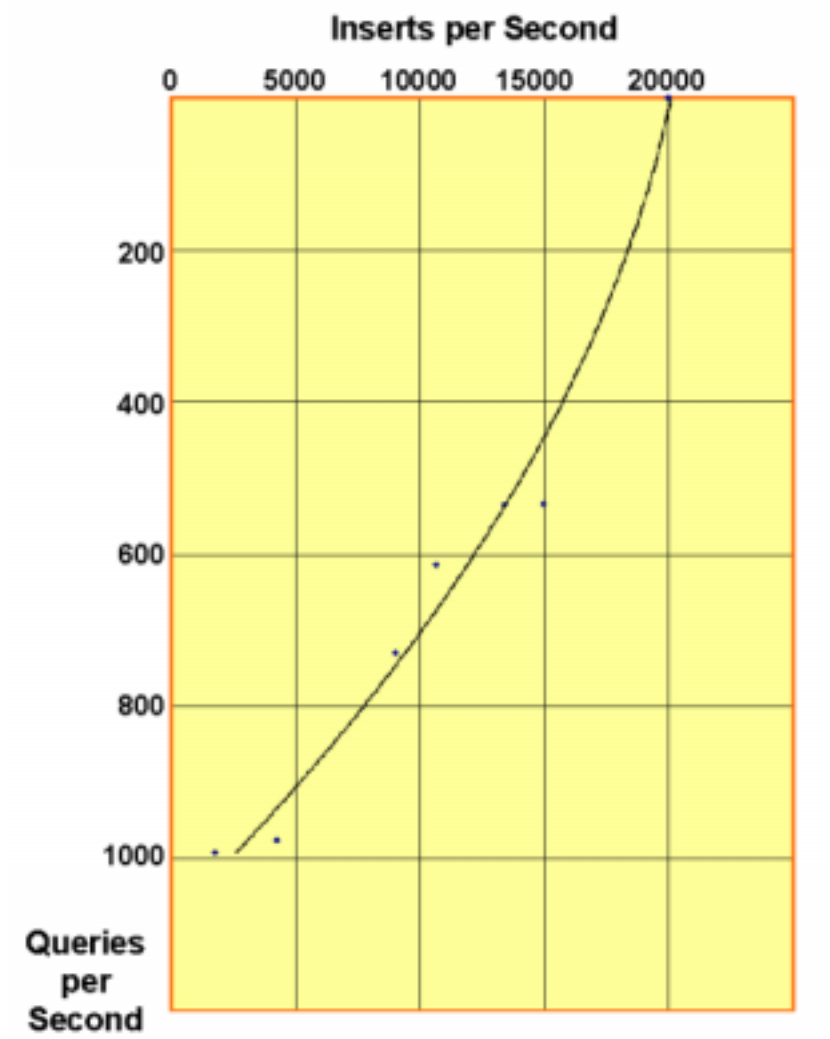

Figure 2 **Phase 3 results**

4. A closer look at the Real-Time ODS Manager

Underlying architectural model

The Real-Time ODS Manager supports both Oracle9*i* and Oracle8*i* in a *TruCluster* environment. Work is underway to provide comparable capabilities on HP-UX. Version 1 of the ODS Manager has been tested and certified with the IONA application server.

Application program interfaces

A client application has the ability to use C++ to interface to the ODS Manager through standard CORBA techniques using the Interface Definition Language (IDL) provided with the installation kit. As an alternative, a Java client can interface to the ODS Manager using the Enterprise Java Bean provided with the installation kit which also works through the standard CORBA mechanism. Or, for performance comparable to C++, the Java client can interface directly to the CORBA mechanism.

Basic functionality

Real-Time ODS Manager Version 1 functionality supports simultaneous processing of the three major classes of transactions that comprise a mixed-workload environment:

- Insert of new records at high speed into the ODS
- Online transaction processing and the ability to do reads, writes, updates, deletes, selects, joins, and other actions against the database
- Traditional decision support queries

Sustained level of transactional service capability

There is a sustained level of transactional service capability associated with the high-speed insert technique. A patent-pending methodology makes it possible to sustain a rate of inserts and never to drop below the level required to achieve overall business goals.

Transaction mapping

Because the ODS Manager must prioritize and control transactions against the ODS database, these transactions must be identified and assigned priority. This would usually be the responsibility of the database administrator (DBA). The transaction identities and prioritization are stored in one or more definition, or "map" files, which contain the following data:

- The definition of the transactions that is used to generate the SQL statement
- The class and type of the transaction, for example, it might be an insert, query, or an OLTP transaction
- The priority of the transaction in the form of a time delay

In regard to the time delay, the current prioritization process associates a fixed, constant time deferral to specific identified transactions that are of lower priority. Inserts always have top priority; then the prioritization scheme for the class of transactions that are queries can be time managed.

The transaction definition and mapping file simplifies the process of transferring the actual data and parameters over the network using client-server communication. The SQL calls are actually formed by the ODS Manager. An application would bind to the ODS Manager and communicate with it. The ODS Manager would encapsulate the Oracle database. Essentially, that is how a higher level of performance is achieved and the balanced mixed workload is ensured.

Unique patent-pending methodology

HP has unique Intellectual Property and a patent pending on this methodology. There are very particular requirements for an ODS, one of which is the ability to insert records at high speed into a database while supporting real-time querying of that information.

As shown in figure 3, this is accomplished with a high-speed insert table with a single index that is partitioned in range.

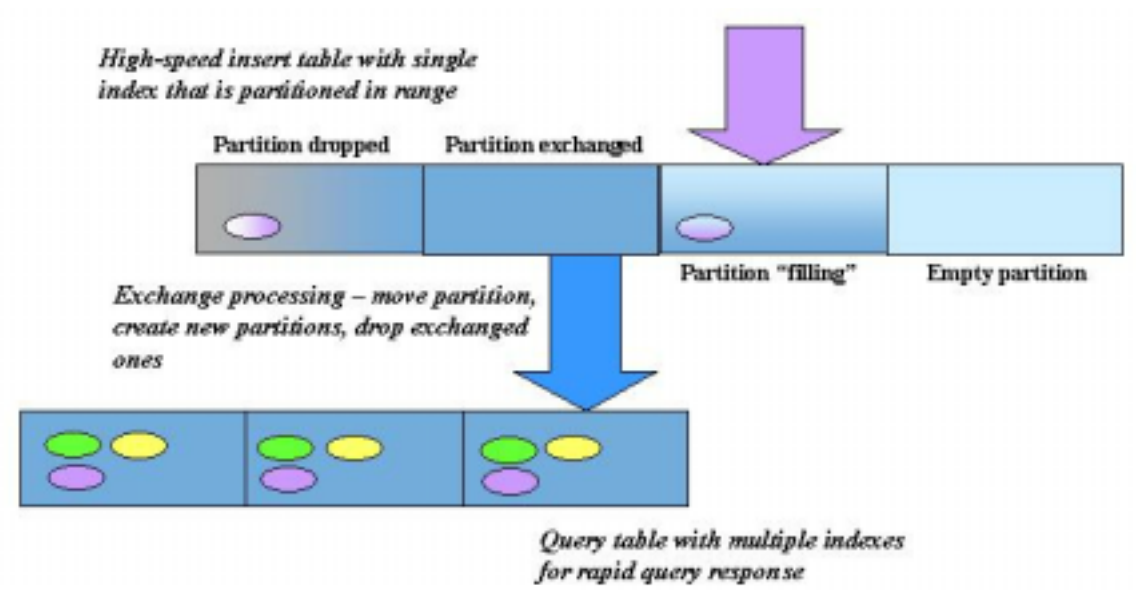

Figure 3 Exchange processing

As transactions arrive – hundreds of millions of telecommunication call data records per day -- tens of thousands per second, for example -- the ODS Manager fills one of these partitions.

The ODS Manager enables the user to specify how often the database exchange process takes place. The high-speed insert table can hold an unlimited number of records. The exchange process empties the table periodically based on the “tunable exchange criteria” set by the user.

When the time limit specified in the tunable exchange criteria is reached, that partition is re-indexed and exchanged into the online data store as “historical” information. These records are re-indexed to provide the level of efficiency required to rapidly query and respond to decision-support and OLTP requests. Then that partition is dropped using a slide window in the “partition exchange” process.

This allows partitions to be filled and then subsequently exchanged using a partition exchange technique. When a partition is dropped, there is always an empty buffer for the next level of in-coming records.

5. Partnerships

Core application services and integration technologies

A full set of tools and applications are required to build a real-time solution. HP has a substantial number of cooperating independent software vendors. They provide “best of breed” capabilities in rules management, message transformation, business process workflow, and ETL (which is the ability to do data extraction, translation and loading). These core application services and application integration technologies are

fundamental to a real-time integration platform and the underlying architecture upon which to build real-time solutions.

Real-time integration partners

Our real-time integration partners are the premier suppliers of these core application services. Except where noted, their software runs on *TruCluster* Server and HP-UX. systems. The full set of application functionality is available on these platforms as shown in Table 1 below.

Functionality	Partner
Rules management	HNC Software
Message transformation	Mercator
Application server	BEA, IONA, Oracle
Application adaptor	Actional
Customer data integration	Acxiom
Data cleansing and de-duplication	Ascential, Trillium
Business process workflow	IONA, SeeBeyond, webMethods
Persistent routing	Talarian
ETL software	Ascential, Informatica
Data mining	SAS Institute
Risk management	SAS Institute
OLAP	MicroStrategy*
Query broadcasting	MicroStrategy*
Campaign management	Unica, Protagona*

*Windows NT only

Table 1 Real-time Integration Partners

Real-time solutions

i-flex
i-flex solutions limited is a world leader in IT solutions for the financial services industry. i-flex enables financial institutions to cut costs, respond rapidly to market needs, enhance customer service levels, and mitigate risk through its suite of proven products and services.

i-flex's flagship product, FLEXCUBE, has adopted the HP Real-Time ODS Manager to provide customers real-time analytic capabilities in the FLEXCUBE Information Center solution. FLEXCUBE is used by more than 100 customers worldwide, and has been ranked among the world's top-two largest selling wholesale back-office banking solutions. Together, HP and i-flex deliver unique real-time value with the ability to:

- Incorporate real-time data from OLTP solutions into analyses
- Deliver real-time alerts to business processes on any communication channel
- Communicate with customers based on knowledge across all touch points

Repeatable solutions

An important note about all real-time solutions based on Oracle databases is that they are repeatable solutions. This means that they lower the cost and accelerate the speed of solution implementation for the end customer.

6. Is an ODS right for you?

You can benefit from an ODS if you need a mixed-workload environment with the continuous insertion of records and the simultaneous processing of queries and transactions in support of real-time decision-making and business processes.

Here are some specific examples by industry of where an ODS is applicable.

- Call detail records (CDRs) in the telecommunications industry
- Point-of-sale records in a retail store environment
- ATM transactions in a consumer banking environment
- Click-stream data information in a Web-enabled application environment

The common denominator in all these examples is that information must be immediately available in real time to be queried or updated and a sustained level of services is required by the business.

Implementing an ODS with HP

Once you have determined that you need access to real-time information and that an ODS is the technology you want to implement, HP can assist you in implementing your real-time solution in a UNIX-based Oracle environment. To best determine the appropriate solution platform to meet your needs, you should ask yourselves certain qualifying questions based on three criteria.

Database strategy

First, what is your organization's database strategy? Is a strategic decision in place to use Oracle software and implement Oracle databases at your company or organization? Are you open to implementing an Oracle database if one hasn't been implemented before?

Real-time business solution

Second, examine the real-time business solution. For instance, your choice of a specific business application for real-time telecommunications fraud management or consumer banking and ATM transactions would lead to a particular platform decision.

Overall IT environment

Third, look at your overall IT environment. What is your mix of platforms and operating systems? What environment is your staff able to operate and support? Are you looking to out-source? Is there a preferred application integration environment; a set of middleware that you use in terms of either standardized application servers or business process and workflow or rules, engines, and things of that nature?

ODS design and implementation services

If your assessment results in a decision to build a UNIX-based ODS, our ODS Design and Implementation Services are ready today to design and build an operational data store for Oracle8*i* or Oracle9*i* databases using Real-Time ODS capabilities. The ODS design services will specify the database schema, the operational parameters to achieve the required performance levels for a specific real-time business workload, and the processes to load and update the ODS. The assessment includes the design information used to make tradeoffs in schema and ODS operational functions to meet the business needs and technical requirements of the applications interfacing to the ODS.

Knightsbridge

Knightsbridge, specialists in high-volume data management, have joined with HP to deliver high-performance data solutions that use HP Real-Time ODS capabilities for Oracle databases. Knightsbridge can help you transform vast data assets into meaningful real-time information. These solutions drive the real-time enterprise—a smarter, faster, and more flexible global business, an enterprise that can use its data for competitive advantage.

7. Conclusion – real-time solutions

A real-time information environment can provide a competitive advantage. A company or organization can achieve that advantage by using unique software developed by HP that combines a high-performance operational data store built with Oracle databases and enterprise application integration technology.

Real-Time Solutions from HP enable a company or organization to recognize patterns and act appropriately when they occur. For the public sector, that ability has never been more critical. And the same is true for companies that compete in information-intensive industries such as financial services, telecommunications, retail, and utilities.

Build UNIX-based real-time solutions today

At the heart of this Real-Time Solutions for Oracle databases initiative is the Real-Time ODS Manager. It gives solution providers the ability to build UNIX-based real-time solutions today.

The ODS Manager is licensable software incorporating patent-pending methodologies that provides the transaction management, database management, and application interfaces to achieve high-performance service levels for transactions. It is both a design methodology for building an ODS and software to manage the techniques employed. What's more, HP's software and database techniques are fully portable to other operating systems

With the ODS Manager, integrators can design operational data stores that support complex, high-volume mixed workloads with a sustained minimum level of service performance. This capability is the core requirement for making business decisions in real-time and is currently not available in a standard database. Using HP cluster systems and the ODS Manager with Oracle*9i*, HP has implemented an ODS that can sustain tens of thousands of inserts per second while supporting hundreds of queries

per second against the same data. Real-Time Solutions for Oracle databases with the Real-Time ODS Manager deliver real benefits:

- Real-time decisions based on up-to-the-second information
- A real-time, integrated view of enterprise data
- Closed-loop business processes
- Real-time management of growing volumes of information

At HP, we invent IT solutions to master the information environment. Our innovations, developed through close industry partnerships, are ready for you to incorporate into an always-on enterprise infrastructure. This is an infrastructure that you can design effectively, manage seamlessly, optimize dynamically, evolve intelligently, and – with Real-Time Solutions for Oracle databases -- leverage operationally to meet your challenges head on.

For more information, go to www.hp.com