

# The Zero Latency Enterprise

Dave Liles

Compaq Computer Corporation  
1295 Stonegate Rd  
Algonquin, IL  
USA  
Dave.Liles@Compaq.com

## Abstract

Today's Internet economy has primed customers to expect immediate access and immediate results. But instant results are difficult to achieve when your customers want up-to-the-minute information about themselves: account balances, detailed transactional histories across all products, immediate problem resolution, and recommended advice on future purchases. Their expectation is that they will get the same interaction with your company whether they communicate with you via telephone, the web, a kiosk, or e-mail. To retain customers and make better decisions, businesses must step up the pace. Customers and systems need access to events – like orders, shipments, and payments – the moment they occur. They need to be able to act on events automatically, in real time. The prescription: a “Zero Latency Enterprise”.

## 1. The ZLE Challenge

From an IT perspective, the challenge is twofold: integrating customer and product information scattered in disparate systems throughout the enterprise, and acting on it in real time.

The first challenge, integrating customer and product

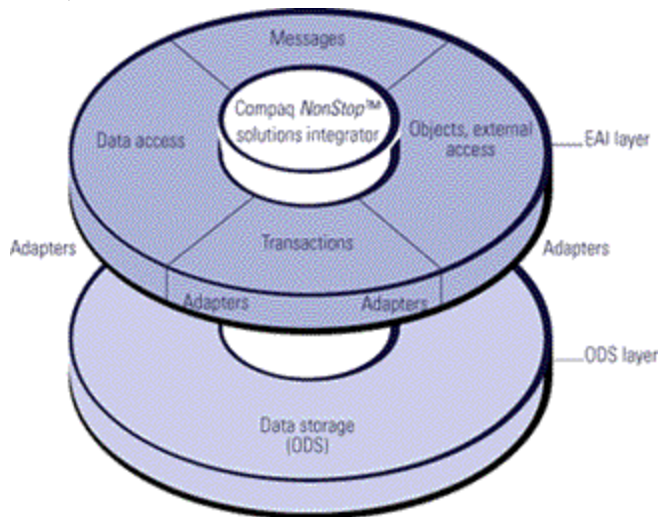
---

*Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the VLDB copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Very Large Data Base Endowment. To copy otherwise, or to republish, requires a fee and/or special permission from the Endowment.*

**Proceedings of the 26<sup>th</sup> VLDB Conference, Cairo, Egypt, 2000**

information scattered in disparate systems throughout the enterprise, focuses on enterprise application integration (EAI). EAI instantaneously transform decisions and information into operational changes by “pushing” information to where it is needed. The second challenge, instant access, analysis, and reaction to information, has been addressed traditionally by using business intelligence technologies, such as data marts, data warehouses, and operational data stores (ODS). Of particular need is a data-centric ODS that is a central, operations-focused data store fed in near real time by all of the other databases in the enterprise. This ODS could provide a central repository from which users and applications “pull” information as needed.

The problems with these traditional approaches is that they don't address the challenge of minimizing the lag, or information float, between when the data is captured in one place and when it can be used somewhere else. Each approach has its limitations; most organizations require a combination of the two – a real time ODS combined with EAI.



However, combining the two approaches has been considered impractical because of technology limitations in providing a single zero latency solution that could:

- Provide 24\*366 service availability
- Loading and updating massive event volumes without disrupting the service
- Extracting data from transactional systems without impacting their performance
- Ensuring consistency of data collected from multiple fragmented source systems
- Scalability in terms of supporting a large, active user base executing a wide variety of both response time sensitive transactions as well as large scale queries
- Very large database manageability

## **2. World's Largest Zero Latency Enterprise Proof-point**

### **2.1 The customer challenge**

A large telecommunications customer approached Compaq with this exact challenge. They asked Compaq to demonstrate a Zero Latency Solution that would provide real time call and network utilization management, enterprise application integration of event data, and real time customer relationship management. They had a 5 terabyte database which was expected to increase in size to 25 terabytes. They wanted to make sure that their platform of choice could handle the expected growth in terms of database size as well as users. They also wanted to mirror the data for reliability.

### **2.2 Compaq's response**

Compaq upped the ante. Compaq decided to build a telecommunication proof-point that would handle the combined business support functions for the equivalent of *the five largest telecommunications companies* in the world! That is, a proof-point that would handle:


- Over 1.2 billion call detail records (CDRs) per day, over 12,000 per second
- An 90 day ODS, over 100 billion events
- Over 45 terabytes of raw data, 90 terabytes mirrored, 111 terabytes of disk storage
- A customer care call center of over 40,000 agents
- Over 20 million customers

### **2.3 The Zero Latency Applications**

Applications that use the ODS include service business processes for real time surveillance of called numbers, customer service inquiries by 40,000 agents, EAI integration processes providing downstream data mart feeds for credit verification, large adhoc business intelligence queries, customer usage analysis (data mining), and system management services via statistics displays.

### **2.4 The Zero Latency Platform**

The center of the Zero Latency Enterprise Solution runs on a Compaq *NonStop Ô Himalaya* Server with 128 processors, 256 GBs of memory, and 111 terabytes of disk storage. An Compaq *AlphaServer GS 140* with 1 terabytes of disk storage was utilized for data warehousing. A combination of multiple Compaq *Proliant* servers were used as simulation drivers, data mining servers, and statistical display servers.



## 2.5 The Zero Latency Solution Integrator

At the center of the demonstration is Compaq's *NonStop* solutions integrator, based on BEA, Java based Blaze (for dynamic rules-based decision making), and CORBA technologies. Message transport services are provided by is Compaq's *NonStop* Tuxedo transaction manager and Compaq's *NonStop* DOM CORBA. Application adapters provide connectivity to outside servers such as Oracle Data Mart Suite, a data mining cluster, and other applications such as SAP. The ODS is based on Compaq *NonStop* SQL database. Large queries against the ODS are performed with Microstrategy's DSS Web tool.

## 3. ZLE Proof-point Results

The proof-point implementation retains over 100 billion CDRs, a 90 day history. Each day, 1.2 billion CDR's are deleted. At peak operation, it accepts up to 50,000 random transactional inserts per second through the CORBA-based application and handles more than 100,000 customer service agents performing more than 3,000 transactions per second in the TUXEDO environment. It also processes 100 customer profile updates per second, physical de-fragmentation of CDR tables, trickle feed to the data mart, and batch extracts to the mining cluster and the massive ad hoc parallel query.

## 4. Conclusion

Compaq's proof-point implementation has exceeded every industry specification and expectation and has confirmed Compaq's revolutionary new concept for

achieving zero latency. The Compaq *NonStop* solutions integrator concept has been validated by customers to serve as a dynamic integrator for applications and a central database for caching data, as well as providing zero latency throughout the enterprise. With our solution, customer profiles and marketing models can be generated dynamically in real time, with up-to-the-second data, in a fraction of the time required by conventional solutions.

This zero latency solution was made possible because the Compaq's *NonStop* Himalaya ODS provides:

- Unprecedented scalability – the shared nothing hardware and software architecture is unique in it's ability to scale incrementally and virtually unlimited in size and bandwidth
- Mixed workload capabilities – the unique capabilities of Compaq's *NonStop* SQL DBMS in combination with Compaq's *NonStop* Software provides for the *simultaneous* execution of demanding transactions, queries, event capture, and application integration
- World renowned availability – patented hardware and software fault tolerance
- Very large DMBS support - *NonStop* SQL DBMS industry leading support for very large tables, online database manageability, and parallel and granular operations