

An Oracle White Paper
January 2011

Oracle TimesTen In-Memory Database For The Communications Industry

Executive Overview	1
Introduction	1
Data Management Challenges	3
Low Latency Data Access.....	3
Carrier-Grade Availability.....	3
Scalability.....	3
Support for Industry Standards	3
Ease of Embedding.....	3
Addressing The Challenges.....	4
Microsecond Response Times.....	4
Durability	4
Data Replication.....	5
Standard Interfaces.....	5
Industry Use Cases.....	6
Wireless Services	6
Voice Over IP	6
Prepaid.....	7
Interactive Applications	8
Conclusion	8

Executive Overview

Communications applications operate under stringent requirements for always-on, fault-tolerant processes. To compete, these applications must be built to ensure minimal latency associated with the retrieval of mission-critical data, provide robustness to ensure non-stop operation and deliver scalability to handle a growing number of concurrent users. This white paper shows how Oracle TimesTen In-Memory Database, which has been widely used in the communications industry for more than a decade, meets these needs and more, providing a strong data management platform for successful operations today and in the future.

Introduction

The rapidly evolving communications industry is going through a fundamental transformation in technology and operational resource utilization. The era of highly customized, proprietary hardware and software is no longer financially viable or desirable as it breeds inherently high costs and lengthy time-to-market—severely limiting service agility. For many years, communications platforms were based on closed proprietary frameworks and data management solutions that emphasized network availability. While the resulting proprietary network infrastructure achieved robustness, it was not designed to flexibly and cost-effectively accommodate new services and technological innovation.

The demand for service agility is driving a fundamental shift in application frameworks. By using the best commercial-off-the-shelf (COTS) solutions, the communications industry is realizing dramatic price and performance benefits and the ability to deploy innovative applications faster. Especially with the growth of electronic commerce and wireless infrastructures leading to expectations of “anywhere, anytime” connectivity, the drivers for choosing an in-memory database for real-time data access are increasingly clear.

To help ensure the highest quality of service, Network Equipment Manufacturers (NEMs), Independent Software Vendors (ISVs) and Communication Service Providers (CSPs) require

extremely fast data access and continuous data availability. In-memory database technology is a vital solution for any application that must supply data instantly and reliably.

Telecommunications and networking applications with high transaction volumes and need for extremely fast response times, are examples of applications particularly well suited to in-memory database technology. In fact, the communications market has used specialized versions of in-memory technology for several years.

Whether running on the enterprise server, in mobile networks, embedded in appliances, or handling complex data that is constantly changing, communications applications need a platform that offers low latency, ensures high availability, and which is built using a scalable infrastructure that allows for rapid growth. To stay ahead of the market and plan for future growth, developers need products that are ready now to integrate into their plans—embeddable products that help provide an agile framework. The products should be easy to use in developing, managing and deploying applications, and should provide a long-term return on investment by offering response time, reliability and scalability that supports high demand and continued growth.

Oracle TimesTen provides the necessary agility for companies developing and deploying communications applications that meet or exceed today's stringent requirements. This paper describes how Oracle TimesTen provides communications application developers superior advantages over building their own — and helps developers deliver greater innovation with shorter time to market.

Data Management Challenges

The twin business challenges of maximizing subscriber retention and monetization of existing and new services have many ramifications for database technology

Low Latency Data Access

Accurate and timely data is essential to the operation of a communications network. Subscribers expect phone calls to go through promptly, data to arrive correctly and other services to be delivered quickly. Tolerance for long wait times is poor and customers always have other options. Solutions that implement policies or have to sift through as many as 10,000 different potential calling plans need access to lots of data very fast. Achieving SLAs of a few milliseconds requires that underlying database transactions complete in microseconds.

Carrier-Grade Availability

Availability is an essential requirement in communications as well as many other industries that operate 24x7 and cannot tolerate service downtime. Solutions that manage session and call information need to ensure that even in the event of a hardware failure, applications will continue uninterrupted. Charging of prepaid accounts, for example, must be not only accurate but must always be available: “your call cannot be completed because we don’t know your account balance” is simply not acceptable.

Scalability

To “death and taxes” can be added a third certainty: growth. While it sometimes stutters in the short term, the long-term trajectory is upwards. And even in the short term, spikes in load occur at different times of the day (when work starts or consumers come home) and different times of the year (holiday periods). Neither performance nor reliability can be allowed to deteriorate as load varies. Any solution must be able scale seamlessly, handling both transient peaks as well as longer-term sustained growth.

Support for Industry Standards

Having identified the right services, development organizations need to build and ship their applications quickly, integrating them into existing infrastructure. Support for the correct industry standards will ensure that developers can use their existing skills and integration into existing applications and tools is much easier, both of which help to shorten time to market.

Ease of Embedding

Fully embedding a commercial database allows NEMs and ISVs to offer complete solutions, delivering services directly and acting as a single point of contact. Since solutions are deployed in many locations, often largely unattended, installation, normal operation, updating and recovery from error conditions all must happen without administrator intervention or an onsite presence.

Addressing The Challenges

Oracle TimesTen In-Memory Database (TimesTen) is a memory-optimized relational database that provides applications with extremely fast response time and very high throughput as required by networking appliances and OSS/BSS solutions including, subscriber and policy management, session and state management, location services, and real-time billing. Deployed in the application tier, TimesTen databases reside entirely in physical memory with persistence to disk storage for recoverability. Applications access the in-memory database using standard SQL interfaces. High availability is provided through real-time transactional replication.

TimesTen can be deployed as either a fully stand-alone in-memory relational database with full persistence and recoverability, or as a database cache integrated with the Oracle infrastructure software stack including Oracle Database, Oracle Real Application Clusters, Oracle Enterprise Manager, and components of Oracle Fusion Middleware.

Microsecond Response Times

TimesTen delivers real-time performance by changing the assumptions around where data resides at runtime. By managing data in memory, and optimizing data structures and access algorithms accordingly, database operations execute with maximum efficiency, achieving dramatic gains in responsiveness and throughput, even compared to a fully cached disk-based RDBMS. TimesTen is designed as an embeddable database within an application to further improve performance of database operations by eliminating inter-process communication and network overheads.

Real-time data management has two performance dimensions – response time and throughput. With TimesTen, a transaction that reads a database record can take less than 5 microseconds (measured on Oracle Enterprise Linux running AMD Opteron 1.8GHz processor). Consequently, throughput is measured in tens to hundreds of thousands of transactions per second, running on commodity hardware, or Oracle Sun Netra carrier-grade servers.

The extremely low latency and predictable response time enables TimesTen to provide real-time access to data for highly time sensitive operations such as real-time call processing and fast-growing call volumes for subscriber bases ranging from thousands to millions.

Durability

There's a natural reaction when hearing about in-memory database technology. It is obvious to most that it will be faster than other approaches, just as RAM is clearly faster than disk. But what happens to all that data when the server goes down?

While TimesTen uses RAM to store the database at run time, the entire database and the transaction logs are persisted to disk in the form of log files and checkpoint files. If the server goes down, the database is recovered from disk. TimesTen supports standard transaction semantics as expected in conventional disk-based relational database systems.

The logging and check-pointing mechanisms are highly flexible, allowing the application developer to tune them as needed. For the utmost in durability, fully synchronous logging to disk or replication to memory on another system ensure there is no single point of failure, while asynchronous mechanisms provide a high degree of durability with much higher performance. Customers have chosen both mechanisms, depending on need.

Data Replication

The TimesTen transaction-log based replication scheme enables high efficiency and low overhead real-time transactional data replication between TimesTen databases. Asynchronous replication provides maximum performance, and the application is completely decoupled from the receipt process of the replicated elements on the subscriber. For applications where transactional integrity is of the highest importance on fail over, synchronous replication provides higher level of confidence for data consistency between the active and standby database; the application is blocked until the transaction has been both received and committed on the standby database.

The flexible deployment architecture supports a range of configuration options over LAN and WAN for active-standby, active-active and N-way replication. The standby database is always available for reads; more read capacity can be provided by configuring additional read-only subscribers.

TimesTen replication is therefore the capability at the heart of addressing several of the challenges outlined in the previous section:

- High availability through integration with Oracle Clusterware, as well as support for online upgrades without downtime.
- Geographic redundancy and locality of data by enabling WAN-based as well as LAN-based replication.
- Scaling through seamless addition of both read-only and read-write capacity.

Standard Interfaces

Applications access the TimesTen database using interfaces such as ODBC, JDBC, OCI, Pro*C/C++ and .Net programming APIs. Unlike many other database systems, where ODBC or JDBC API support may be much slower than the proprietary interface, ODBC and JDBC are native TimesTen interfaces that operate directly with the database engine. TimesTen supports versions of these APIs that are both fully compliant with the standards and tuned for maximum performance in the TimesTen environment.

TimesTen also support the Oracle Call Interface (OCI) and the Pro*C/C++ pre-compiler for TimesTen functionality. TimesTen OCI and Pro*C support enables you to run many existing OCI applications with TimesTen in direct mode or client/server mode. . Furthermore, those applications that need access to TimesTen from the MS .Net framework can do so using the Oracle ODP.NET data provider.

Industry Use Cases

Oracle TimesTen has been a part of mission-critical infrastructure in many different parts of the communications industry for over a decade. In this section we will examine some different use cases to show how industry challenges are being met and where TimesTen is being deployed.

Wireless Services

A breakthrough in network infrastructure applications will be enabled by the adoption of a platform that enables application developers to focus on differentiated services, relieving them of the needs to worry about data management issues. The role of the subscriber profile is key to extending and blending existing offerings in combinations of data, video, voice, and content assets. CSP's efforts to manage, enrich, capitalize on, and provision user information across subscriber databases are critical to fuel applications-led service initiatives.

Subscriber and policy management systems act in concert with network infrastructures, and most investments will link to other next-generation infrastructures as well as legacy infrastructure. For example, IP-based core systems will drive steps toward the need for securitized user profiles. Adjunct systems such as charging will also be associated with subscriber management investments as the concept of what the user can access relies on the payment agreements.

The rapid growth of 3G data services has put unprecedented demands on the mobile network infrastructure, the with the advent of 4G and LTE and predicted to fuel further growth and new services in areas such as machine to machine (M2M) connected services, wireless solutions need to meet critical availability, reliability and response time requirements.

Oracle TimesTen powers many highly scalable, reliable and modern subscriber management solutions, from home location registers to identity administration and home subscriber servers.

Voice Over IP

With VoIP becoming the preferred means to extend telephony services, its underlying network elements must deliver the same quality of service associated with traditional telephone services. Many VoIP network elements were developed a number of years ago, but most have not been engineered to meet the requirements of large enterprises and communication providers who plan to use it as the backbone of their offerings, supporting millions of subscribers. As a result, Telecommunications Equipment Manufacturers (TEMs) and Independent Software Vendors (ISVs) are reassessing their current solutions to ensure they can deliver speed of innovation, scalability, and reliability. The foundation for any VoIP network element necessitates a new breed of infrastructure software purpose-built to meet such requirements.

Data availability and response time are crucial to effective operation of the VoIP network. The availability of the most up to date call routing information ensure timely call connection and use of the most efficient or least cost routes across peer networks. High available of cached cal session state ensures active calls can be maintained and appropriately billing in the event of a server outage.

Oracle TimesTen can be found inside key VoIP network elements including:

- IP PBX and IP Centrex/Hosted PBX
- Softswitch/MGCF and Gateways
- SIP Servers and Proxy Servers
- Service Delivery Platforms
- Application Servers
- Network and Quality Management Systems
- IN SCP
- Session Border Controllers

Prepaid

Prepaid charging was once exclusively for subscribers considered at risk of defaulting on payments. Around the world, prepaid charging is now a growth segment targeting teens, students, and seniors. For this segment, communication service providers need more flexibility in managing subscriber balances and how much money is spent, which is driving the following requirements:

- **Offload existing intelligent network (IN) solutions**—enabling adjunct prepaid applications to offer innovative discounting, service bundles, and promotions.
- **Tie prepaid accounts to high-paying postpaid accounts**—a postpaid parent creates a prepaid account for their teenager with spending restrictions and automatic recharge from their account.
- **Provide postpaid subscribers the ability to trial new services on a prepaid basis**—a postpaid subscriber tries 3G data access using prepaid to eliminate the risk of spending more than desired or having a monthly commitment.
- **Terminate service instantaneously to reduce lost revenue**—notifying subscribers when their balance is depleted and giving them options to recharge.

Oracle TimesTen is used at the core of prepaid charging applications that enable communication service providers to

- Scalable throughput to rate CDRs and IPDRs
- Deliver real-time balance updates, account replenishing, and advice of charge (AoC)
- Instantaneously terminate services when needed
- Integrate with billing solutions and other applications
- Support usage- and content-based rating requirements
- Support multiple services on a single platform

Oracle TimesTen manages all performance- and time-sensitive information used in prepaid charging:

- Balance and rating
- Authentication and fraud policies
- Subscriber and services profiles

- Discounts, promotions and usage counters

Oracle TimesTen In-Memory Database provides the reliability, scalability, and instant responsiveness to accommodate the increasing volume of network events handled by prepaid charging applications. Oracle TimesTen can eliminate any single point of failure by distributing workloads across multiple platforms for load balancing and fail-over protection. This reliability is achieved without sacrificing fast processing and rating of CDRs and IPDRs or the ability to manage complex subscriber profiles. Moreover, Oracle TimesTen is based on industry-standard interfaces, providing the flexibility to evolve and to easily share data with other applications and OSS/BSS systems.

Interactive Applications

In contrast with machine-to-machine communications, applications involving human interaction can seem positively pedestrian. Response time requirements are typically in the range of 1-2 seconds, rather than milliseconds – up to 1,000 times slower. So at first glance it seems somewhat surprising that in-memory database technology is appropriate here. And yet, there are significant benefits to be had from accelerating this class of application.

Firstly, both customers and staff appreciate and expect predictable response times. As we have seen, sometimes loads can spike during peak periods and this does slow response time. So a perfectly acceptable 1-2 second delay becomes a 10 second delay. Multiply this by the larger-than-usual number of interactions, and that's a recipe for poor customer service that nobody wants.

But it's not just about avoiding negatives; there are positive possibilities as well. Customers accessing websites for self-service applications or similar, will notice and appreciate a much faster response time. This can make web portals an attractive way for them to interact with service providers, both reducing costs and increasing the potential to offer them more advanced services (with the attendant possibility of incremental revenue). In-memory technology is already used on many retail websites to provide personalization and improve the overall experience; the opportunity to differentiate other kinds of online presence is not yet widely taken.

Finally, in call centers, service reps often spend time waiting for basic account and profile information. Accessing this much more quickly does more than speed up the transaction. It enables reps to get complete information, identifying up-sell and cross-sell opportunities, again generating incremental revenue.

Conclusion

Particularly for applications in the call path, the communications environment is a challenging one. The twin needs for low latency and carrier grade reliability are not easily met, particularly during times of peak load. Oracle TimesTen is proven in the communications industry, embedded in solutions used by hundreds of millions of consumers worldwide. Leading communications companies today rely upon Oracle TimesTen to deliver both low latency and carrier grade reliability in their most demanding environments today, with scalability built-in to handle the growth they expect and plan for.



White Paper Title
[Month] 2011
Author: [OPTIONAL]
Contributing Authors: [OPTIONAL]

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200

oracle.com



Oracle is committed to developing practices and products that help protect the environment

Copyright © 2011, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. UNIX is a registered trademark licensed through X/Open Company, Ltd. 1010

Hardware and Software, Engineered to Work Together