

AMD64 Multi-Core Computing Platforms and solidDB enable consolidated approach to subscriber databases and related applications

Executive Summary

With solidDB and AMD64 Multi-Core Processors, system architects can consider a new, consolidated approach to subscriber databases and related applications. In this new approach, the applications are run in the same multi-core/multi-processor system, together with the database server.

With the consolidated approach, benefits can include greater performance, higher processor utilization, and improved load balancing and lower administration costs. In the consolidated approach, applications communicate with the dataset via the high bandwidth and low-latency Direct Connect Architecture, to decrease external network traffic and improve efficiency. AMD Virtualization™ technology may further improve performance when virtualization is used to place applications and database in different virtual machines. High availability is achieved with solidDB's support for always-on data access; applications are restarted as they would be if they were on a separate system. While not applicable in all situations, consolidation is an attractive new approach many telecommunications companies should consider.

As shown by the test results, the combination of solidDB and AMD processors with Direct Connect Architecture proves to be an effective solution to be considered for database deployments. This is a key point for the requirements of the new dynamic (always changing) style databases that are being driven by new types of services; IPTV, Broadband, MobileTV, VoD and other flexible offerings that can be turned on and off by the subscriber in a heartbeat.

This white paper briefly reviews subscriber databases, and then discusses AMD64 technology, the solidDB database, and benchmarking approaches. The white paper

concludes with the results of a standard TMI benchmark test, comparing client/server and configurations.

Subscriber Databases in Telecommunications

The central subscriber database of the mobile communication network is called Home Location Register (HLR). It contains details of mobile phone subscribers of a specific operator. An HLR exists for each mobile telephone network and operator. It may reside in one computing node or it may be distributed over several nodes for scalability and reliability reasons. An HLR entry corresponds to a SIM card issued by the operator. The unique identifier of a SIM card (IMSI) is indeed a primary key of each HLR record. Furthermore, an HLR record contains the primary and secondary telephone numbers of a subscriber, a list of services the subscriber is allowed to use, GPRS settings, the current location of the subscriber and call divert settings. The HLR receives and processes MAP transactions and requests from the Gateway-Mobile-Switching-Center, from the Visible Location Register, the Short Message Service Center, and the Voice Mail System. If the HLR fails, then the network is effectively disabled.

The comparable function in the IP-Multimedia-Subsystem is called Home-Subscriber-Server (HSS). Like the HLR in mobile telephone networks, the HSS is the master user data base supporting the functions that actually handle the calls/sessions. The HSS stores subscriber related information, stores user profiles, performs authentication of users and provides information about the physical location of the subscriber. The HSS connects to the Serving-Call/Session-Control-Function (S-CSCF), Application Server (AS) and IP-Multimedia-Service-Switching-Function (IM-SSF).

Issues and challenges with subscriber databases

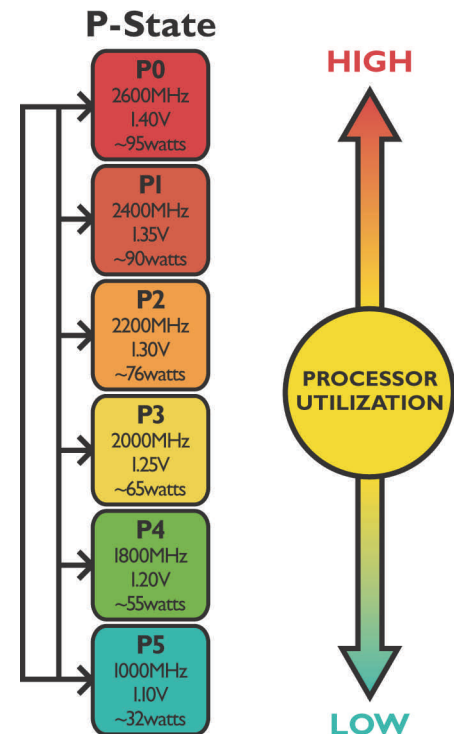
The subscriber database designer faces four major challenges: performance, reliability, scalability and versatility. Depending on network size and number of subscribers, performance requirements may vary greatly. Performance of subscriber databases is measured by the number of lookups per time unit. Since the subscriber database is a centralized database, reliability measured in seconds of downtime per year, is essential. Deployment of networks happens gradually. Operators are interested in incremental infrastructure investments to optimize the return on investment. Therefore, scalability becomes an important criterion too. Last, but not least, the ability to handle various forms of networks imposes a challenge to the system designer.

AMD64 based Computing Platforms for Telecommunications

From a hardware point of view, key equipment within Next-Generation-Networks, is based on server computing technology. Beyond standard server products for Next-Generation-Networks, AMD Opteron™ processors are also available in open standard chassis platforms like AdvancedTCA (ATCA), BladeCenter and CompactPCI. These platforms are building blocks for solutions that provide the longevity and reliability required by service providers in communication networks such as IP Multimedia Subsystem (IMS), Voice over IP (VoIP), Global System for Mobile Communications (GSM), and Code Division Multiple Access (CDMA) wireless networks. At this level, AMD64 technology provides several unique customer value propositions:

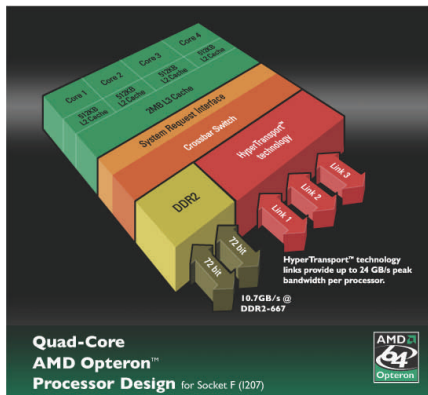
Industry-Leading Performance-Per-Watt

With the equipment changes and upgrades to accommodate new services and additional throughput, the costs from power consumption and air conditioning are becoming a major concern for service providers. AMD PowerNow!™ technology with Optimized Power Management provides power on demand computing, and can reduce CPU power consumption by as much as 43 percent while under an approximate 60 percent load, or up to 75% reduction in CPU power when idle. Adjusting CPU power for the load creates a ripple effect of benefits in the data center that can increase significantly as the number of processors increases. These benefits can include reduction in the amount of heat produced by the processor and other components, and reduction in the amount of electricity required to both run and cool the system. By offering a processor that provides for a dense compute configuration with low overall power consumption, AMD helps make it possible to put more compute capability in data centers and central offices and potentially lower utility costs.



Scalability

Evolutionary deployment of new technology like IMS and 3G must be based on scalable architectures and the long-term protection of past investments. To this end, AMD64 technology offers an innovative solution that allows network operators and service providers to run 32-bit and 64-bit applications simultaneously – operators and service providers can grow into 64-bit computing without sacrificing their existing investments.



Perhaps even more importantly, AMD64 technology was designed from the ground up with multi-core computing in mind, resulting in high performance for multi-threaded and multi-tasking environments. Quad-Core AMD Opteron processors will be socket compatible with today's DDR2-based AMD Opteron processors, offering a seamless upgrade path to quad-core performance as demand increases.

This simplified platform approach to implementing multi-core technology is further enhanced by AMD's revolutionary Direct Connect Architecture, which provides exceptional system performance and multi-processor scalability in up to 8-processor systems.

AMD's Revolutionary Direct Connect Architecture

AMD64 technology with Direct Connect Architecture was designed to improve overall system performance and efficiency by directly connecting the processors, the memory, and the I/O. Directly connecting these elements helps eliminate the front-side bus bottleneck inherent in traditional x86 architectures. On the AMD Opteron processor, the integrated DDR2 DRAM memory controller changes the way the processor accesses main memory, to increase bandwidth and reduce memory latencies. With dual channel RDDR2-667 available on the AMD Opteron 2000 Series and 8000 Series processors and UDDR2-800 available on the AMD Opteron 1000 Series processors, theoretical maximum memory bandwidth (per core) is between 10.6GB/s and 12.8GB/s. Another benefit of integrating the memory controller directly on the processor is that memory bandwidth scales with the addition of processors. This helps provide high-throughput responsiveness and scalability with multi-processor platforms.

AMD64 technology with Direct Connect Architecture features HyperTransport™ technology and provides up to 8.0GB/s bandwidth per HyperTransport link— sufficient bandwidth for interconnect technologies including PCI-Express®, PCI-X®, DDR, InfiniBand, and 10G Ethernet. AMD Opteron processors support up to three HyperTransport links per processor, providing up to 24.0GB/s peak bandwidth per processor.


These processor-level features can provide a tremendous competitive advantage for Home Location Register (HLR) in GSM / CDMA Networks and Home Subscriber Server (HSS) in IMS networks. AMD's Direct Connect Architecture allows each processor in the multi-processor compute node to access the memory of the other processors at low latency with high bandwidth. The up to 256 terabytes of memory address space of AMD

Opteron processors help the performance of applications when large (or many) datasets are held in memory, as is the case with subscriber data bases in HLR and HSS.

AMD Virtualization™ Technology (AMD-V™)

Second-Generation AMD Opteron processors with DDR2 memory are designed to help end the “one server, one application” problem. With its hardware-assisted virtualization using AMD Virtualization™ technology (AMD-V™), Second-Generation AMD Opteron processors can help streamline the Telecom infrastructure to achieve higher levels of efficiency and utilization.

Regardless of the specific network implementation chosen by a particular service provider, AMD-V in conjunction with a hypervisor helps enable a flexible, one-to-many relationship between hardware, operating systems, and applications, and contributes to functional flexibility, implementation, and distribution for network equipment vendors.



solidDB Database Solutions for Telecommunications

Solid Information Technology is the leading provider of fast, always-on and affordable database solutions. Solid serves enterprise and OEM/embedded customers with a combination of open source and proprietary database products. Hundreds of leading companies in multiple industries rely on Solid for their mission-critical applications. There are more than 3,000,000 deployments of Solid's database worldwide in telecommunications networks, enterprise applications, and embedded software and systems.

solidDB™ delivers standards-based, full-featured relational databases specifically designed for high performance data access and always-on data availability to satisfy real-time demands. It enables you to embed a robust and compact database directly into your solution for extremely fast data access that solidDB makes always available (99.9999% uptime). solidDB has the unique ability to deploy an in-memory database in conjunction with a disk-based database, thus optimizing price/performance from within a single programmatic interface, transparent to the application. This "hybrid" characteristic also applies to advanced functionalities such as data replication and high availability, significantly reducing the complexity, time and cost of developing new applications.

Additionally, because all aspects of solidDB can be controlled and fine-tuned within the application, no DBA is required.

Hundreds of Solid customers across the world have chosen solidDB because it:

(1) Achieves real-time performance: solidDB combines multi-threaded main memory and disk-based database engines to achieve fast transaction speeds and the most effective utilization of system resources within a single database product.

(2) Provides uninterrupted data access: with a unique hot-standby architecture that is transparent to applications, solidDB enables applications to recover from system failure within milliseconds, delivering carrier-grade high availability (99.9999%) that leading carriers need.

(3) Is flexible to adapt to change: solidDB can be embedded and programmatically instructed to initiate additional database instances on demand, move the data closer to users and provide access to transaction-consistent data at any scale and from any location.

(4) Is affordable: In addition to product capabilities that address the most stringent service level agreements, solidDB provides attractive total-cost-of-ownership, rapid development time and low operational costs.

Designed for advanced, 64-bit multi-core architectures

solidDB is architected to exploit advanced systems based on AMD64 technology: multiple cores, high bandwidth interconnect and large memory capacity. These advanced systems contain from one to as many as eight processors (up to 16 cores). solidDB can support distributed database instances on multiple nodes, as well as single database instances on large multi-processor systems. In either case, solidDB leverages the parallelism in multi-core processors, whether existing dual-core designs, the coming quad core designs, and beyond. The combination of solidDB and AMD64 technology deliver the performance, reliability and affordability demanded by mission critical telecommunications applications.

solidDB for MySQL

solidDB for MySQL is a turbo-charged version of MySQL database that leverages solidDB Storage Engine to support heavy transaction volumes and large numbers of concurrent users. By leveraging Solid's proven online transaction processing (OLTP) storage engine,

solidDB for MySQL is able to support large numbers of concurrent client connections and multiprocessor/multi-core hardware configurations.

Performance Benchmarking for Subscriber Database Implementations

For telecommunications systems managers, benchmarking plays an essential role in optimizing performance, capacity planning, and selecting the right hardware and software products. With ever more demanding services getting deployed at an ever increasing pace, the consequences of a bad decision are high: lost revenue, higher churn rates, and delayed adoption of new services. With telecommunications companies shifting to low-cost systems, it's essential to get sizing and performance right.

There are two kinds of benchmarks: special purpose and standardized. A special purpose benchmark is written by the development team to test the speed of a representative sample of the application. This kind of test may provide quite accurate estimates of application speed. However, such benchmarks are costly to write, and delay the development task. Benchmark applications require additional functionality for running multiple clients concurrently, collecting timing information and analyzing the test results. Also, because each test is unique, there is no way to compare results to tests run by others.

Standardized benchmarks may resemble the target application less precisely, but they provide a common way to define and measure test configurations. Over time, the multiple tests run by various parties constitute a database of comparative performance data. Developers can use this information to select a configuration that is likely to be the fastest for their application, without ever having to run a benchmark for themselves, saving both time and money.

Existing database benchmarks that come from the Transaction Processing Performance Council (TPC: see www.tpc.org) center around enterprise applications: order entry, decision support and web server applications. While standardized, these have little relevance to those building real world applications with very different database usage profiles.

The TMI Database Benchmark

TMI is a standardized database benchmark that, unlike the TPC benchmarks, was developed to address demanding telecommunications scenarios. The TMI benchmark simulates an HLR application. Every call to and from a mobile phone involves look ups

against the HLR, making it a central part of many telecommunications applications, and often creating extremely large numbers of short transactions. An HLR database holds subscriber identification information as well as services to which the customer has subscribed, the subscriber's current location, preferences and billing details.

The TMI benchmark is widely used by Network Equipment Providers, as well as software, hardware and processor vendors, to compare performance of different products. TMI originated as a special purpose benchmark used internally by a telecommunications equipment manufacturer. After the specification was published by the University of Helsinki in 2004, Solid Information Technology enhanced the specification and implemented the benchmark. Later that year Solid published the entire TMI benchmark description on its web site and made a benchmark kit available for free download. Recently, Solid contributed the TMI benchmark suite as an open source tool under the GNU General Public License (GPL). A related benchmark, called Network Database Benchmark, published recently by Nokia is also used to evaluate databases for telecommunications applications.

The TMI database schema is a simplified version of a real-world HLR application schema. The transactions that run against the database, creating the load, are taken from real-world HLR applications. This makes TMI the benchmark of great relevance to telecommunications equipment manufacturers and network infrastructure application developers.

Technical details of the TMI benchmark schema and transactions may be found at <http://www.solidtech.com/tmi>.

A benchmark run consists of five stages, as shown in Figure 1 below. The entire run takes the time sufficient to obtain reliable results.




Figure 1 TMI benchmark stages

Step 1 is to create the database schema, and populate the tables with randomly selected data that fits the TMI parameters. Next the RDBMS is allowed to idle for a few minutes so that it can complete any background asynchronous tasks, such as checkpointing. During

the ramp-up time, the TMI transactions are run against the database, but results are not collected. This period warms-up the RDBMS cache (if used) with a random set of records from the database, reflecting the pattern of use. Finally the TMI transactions are run under controlled and timed conditions, and performance data is saved for later analysis. The analysis results in a “Mean Qualified Throughput (MQTh)”, measured in transactions per second, suitable for comparisons between different configurations.




Figure 2 Typical TMI client /server test configuration

The client server configuration is shown in Figure 2. Note, there can be more than one client in a system. In a consolidated configuration, the TMI clients are placed in the test system.

Benchmark configuration

TMI benchmarks are traditionally run with load generators on “client” systems sending simulated HLR transactions across a network to a database on another system. This approach presumes an architecture where applications and the database reside on different systems. To test the benefits of a consolidated system, where both applications and database are on the same system, a new approach was required. In the absence of a benchmark suite specifically designed for consolidated systems, a relatively simple approach involves placing both the TMI load generators and the database being tested on the same system. Network traffic is reduced, as would be the case in a genuine consolidated system, while the load generators essentially “stand in” for the applications. The results can help determine the performance benefits of consolidated systems, but can’t be used for capacity planning purposes beyond the database component.

The data base server under test used the TYAN motherboard version S4881 equipped with eight Dual-Core AMD Opteron Model 865 processors. The system deployed 4 GByte

DDR memory per processor, or 32 GByte DDR memory for the entire data base server. The disk system was attached via a PCI-X SCSI RAID Controller from LSI Logic. The RAID controller deployed 128Mbytes cache. The size of the disk system was 2 Terabytes, although only a small percentage of that space was used.

Software on the data base server consists of the operating system, SUSE Linux Enterprise Server 10, solidDB 6.0 Beta, with in-memory engine (IME) and TMI version: 3.0.6.

The database was populated with 100,000 subscribers, accessed by 32 client processes. Checkpointing and logging were disabled. The transaction mix was set to 80% read accesses and 20% write accesses. The rampup time was limited to 1 minute and sampling time was set to 10 minutes. The short rump-up time was justified by the fact that, in an in-memory database, the database cache (buffered page pool) is not used.

Benchmark Results

Two benchmark runs were made. During the first run the classical client/server setup was tested and during the second run the consolidated setup was tested. Results are shown in Figure 3.

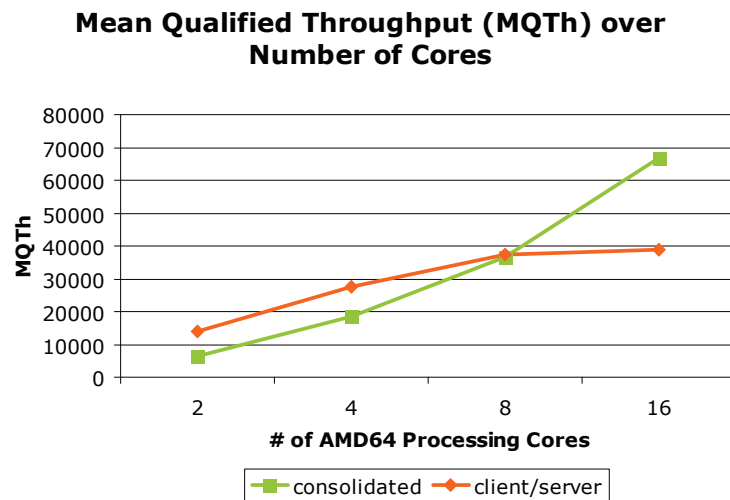


Figure 3 Number of transactions depending on number of cores and approach

The benchmark results for the consolidated configuration indicate that the combination of solidDB and AMD64 technology delivers exceptional performance with large numbers of cores. With a small number of cores operating, the cycles consumed by the load generators outweighed the savings in network overhead associated with a consolidated system.

However with larger numbers of cores, load balancing and network overhead savings outweighed the presence of the load generators on the system, resulting in large economies of scale. In the consolidated tests, a 16-core system outperformed a 2-core system by more than ten-fold, resulting in a scalability factor of 1.25 (with regards to the number of processors).

Conclusion

A new consolidated approach, where applications and database reside on the same system offer communication providers significant performance using cost effective industry standard technology. The TMI benchmark demonstrates Solid's ability to leverage AMD's multi-core technology and AMD's revolutionary Direct Connect Architecture. Consolidated systems based on solidDB and AMD64 have become an attractive option worth considering.

Additional resources



AMD Opteron processors from AMD: <http://www.amd.com/opteron>



solidDB from Solid Information Technology: <http://www.soliddb.com>

TM1 benchmark: <http://www.soliddb.com/tm1>



Advanced Micro Devices
One AMD Place
P.O. Box 3453
Sunnyvale, CA 94088-3453
www.amd.com
www.amd.com/amd64embedded

©Advanced Micro Devices 2007. All rights reserved. AMD, the AMD Arrow logo, AMD Opteron, and combinations thereof, AMD PowerNow!, AMD Virtualization, and AMD-V are trademarks of Advanced Micro Devices, Inc. PCI-X and PCI Express are registered trademarks of PCI_SIG. Other names used in this document are for identification purposes only and may be trademarks of their respective companies.