Intel® Xeon® Processor X5570-based Blade Servers Deliver Unmatched Performance for InterSystems® Caché® Database Benchmark

Abstract

In November of 2009, Intel and HP joined InterSystems® in benchmarking a commercially available Electronic Medical Records (EMR) management system based on InterSystems Caché® database. The benchmark utilized HP ProLiant* server blades based on Intel® Xeon® processor X5570.

The benchmark revealed the InterSystems Caché-based solution achieved unmatched performance and scalability, sustaining an average of 768 episodes/minute (over 11.1 million database accesses/second), which is equivalent to supporting 67,200 concurrent users with a single 16-blade, enterprise-level server system.

Table of Contents

Executive Summary

Database performance is a critical component to massive transactional processing systems, such as Electronic Medical Records (EMR) management applications. Traditionally, costly IT solutions have been required to support scalability of today's large relational database deployments using object-oriented programming technologies.

This benchmark, hosted by HP and supported by InterSystems, Intel, and VMware, reveals the high scalability and performance of InterSystems Caché running on quad-core Intel® Xeon® processors X5570, in both a native configuration and a virtualized configuration. The combination of Caché and the Intel Xeon processor blade-based server achieved performance and scalability results not seen in these configurations before.

Database Performance Challenges in EMR

Database performance for EMR management is a growing concern for hospital IT personnel and EMR application developers. As large hospitals and hospital groups migrate their patient information to EMR, the database must adequately scale and support a massive number of database accesses to allow administration personnel to add, update, and retrieve patient information thousands of times per day, in some installations.

Historically, to achieve necessary performance, deployments relied on in-memory database implementations, which require massive amounts of costly memory. In addition to high deployment costs, the inherent reliability risk with these types of solutions for business-critical applications is their inability to persist data. Today's implementations typically use persistent, relational databases.

As persistent, relational databases supporting modern object-oriented applications scale to handle larger volumes of data, their performance often degrades. This is usually because of the inherent mismatch between object-oriented development technologies and the two-dimensional, rows-and-columns data structures used on disk. The processing overhead required to "map" between the complex data types used by today's sophisticated software applications and a relational database tends to restrict throughput.

Large, monolithic architectures, symmetric multi-processor (SMP) systems, clusters, and grid computing have been implemented to accommodate processing overhead demand to meet required performance objectives with relational database systems and object-oriented applications. However, big machines are very expensive, and large, distributed systems can be difficult to manage, driving up overall total cost of operations.

The Caché/Blade System Combination

InterSystems Caché object database provides an underlying database architecture ideal for applications based on object-oriented technologies. (Caché is seamlessly accessible via ODBC and JDBC as well, so reporting tools that rely on SQL can be used without mapping.) In addition to eliminating the mapping overhead of traditional relational databases, Caché has proven to scale extremely well while delivering outstanding performance for large installations, as shown in this benchmark.

Today's blade servers based on quad-core Intel Xeon processors X5570 offer scalable, high-performance, and cost-effective IT solutions for massive transactional processing requirements. In addition, advanced processor technologies, including virtualization enhancements, power savings, and additional features, further benefit small to large IT deployments.

This benchmark illustrates the unmatched performance and scalability achievable with the combination of the Caché object database design and Intel Xeon processor X5570 in a blade-based server for both native and virtual environments.

Methodology

Several benchmarks delivered performance, scalability, live migration/failover, and encryption burden results:

- A scalability and performance benchmark with up to 12 application servers and up to four benchmark drivers.
- A performance benchmark with a single-server, native configuration with one benchmark driver.
- A performance benchmark with four virtual machines (VMs) and a single, benchmark driver.
- A live migration test from one blade to another of an application server in a virtual environment.
- An encryption impact analysis of database encryption on CPU performance.

Test Drivers

Based on usage patterns observed at a large InterSystems production installation, InterSystems engineers developed task activity scripts to simulate web-based clients (called drivers in these benchmarks). These scripts performed typical EMR tasks. Simulated patient data, also derived from large real-world installations, populated the Caché database.

Test Background

A task creates one or more episodes, which is a patient journey or visit to an inpatient, outpatient, or emergency facility. Each episode requires multiple transactions carried out by application components and results in multiple database accesses.

Database accesses per second and user activities (measured in episodes per minute[1]) correlate to CPU consumption and disk IO activity. The higher the number of database accesses/user tasks the more work the system is doing, resulting in a rise in CPU utilization and disk IO.

During the benchmarks, the number of simulated web clients and application servers varied across the configurations.

- Up to four benchmark drivers were used for scalability testing with 16 blades online.
- Application servers scaled linearly from 1 to 12 servers for scalability testing.
- Only one benchmark driver was used for single-blade native configuration and virtualization tests.

Engineers monitored throughput and various response times to determine if a test run was valid. They recorded several parameters, including sustained and peak rates of database activities in terms of episodes/minute and database accesses/second. In addition, CPU utilization was captured.

Benchmark Architecture

HP hosted this performance benchmark at their PTAC center in the U.S. Tests were conducted in October-November, 2009. The benchmark was completed using HP ProLiant server blades with Intel Xeon processors, a commercially available EMR application based on InterSystems Caché database, and simulated clients. VMware vSphere* and vCenter* were used to test performance and migration/failover in virtual environments.

Table 1 describes the hardware configuration and software used.

InterSystems Caché

InterSystems Caché database is a multi-dimensional, object database. Unlike traditional relational databases, Caché was designed to support today's object-oriented application technologies, with a focus on performance, flexibility, developer support, and scalability.

InterSystems Enterprise Cache Protocol (ECP)

InterSystems Enterprise Cache Protocol (ECP) accelerates transaction activities by caching data locally to each Caché application server. Applications access local cached data first before requesting remote data. Responses from the remote database include both the requested data and the data of the entire block where the data was stored on the remote database. ECP speeds transactions and reduces network traffic between the application and database servers.

Intel Xeon Processor X5570

Built on the Nehalem microarchitecture, the Intel Xeon processor X5570 is a quad-core server processor designed for performance and efficiency with efficient power management to help reduce operating costs. The processor supports Intel® Virtualization Technology,[2] enhancing virtualization performance.

Table 1. Benchmarks Hardware/Software Components

Component	Brand	Configuration
Servers	• HP ProLiant* BL460c G6 server blade • HP c7000 BladeSystem* enclosure • 16 blades in the system	Per-blade configuration: • Dual-socket server board • Intel® Xeon® processor X5570 @ 2.93 GHz • Four cores/processor • Hyper-Threading enabled • (Total of 8 cores and 16 threads per blade) • 48 GB memory/blade
Storage	HP StorageWorks* 8400 Enterprise Virtual Array	8 GB host ports, 128 15k Fibre channel disks
Network	HP Virtual Connect Flex-10 10 GB Ethernet	16x10 Gbps downlinks to blades
Database	Caché® Version 2009.1	Two-node implementation with InterSystems ECP layer on top of the database layer
Test bed	Commercially available Electronic Medical Records (EMR) application	12 application servers (12 blades: dual-socket, four cores per socket with Hyper-Threading enabled: 16 threads per server)
Web server	Apache* 2.2	N/A
Virtual Environment	VMware vSphere* + vCenter*	Single ProLiant blade configured as follows: 3 VMs, 4 virtual processors per, 12 GB memory per VM VMotion* was used to test live migration of an application server
Operating System	Linux* 64-bit	N/A

Scalability Benchmark Configuration Overview

Up to four benchmark drivers simulating web sessions (or users) executed application workflows. Up to 12 application servers running a commercially available EMR application serviced the tasks. The Caché database ran on one node with a second node for failover testing (this node was not used in the scalability benchmark). InterSystems Enterprise Cache Protocol (ECP) layer provided efficient ECP client-side caching for a truly distributed database cache. See Figure 1.

The scalability/performance benchmark ran on a 16-blade system:

- 12 blades provided the application servers
- 1 blade hosted the database application (with a second for failover testing)
- 2 blades were used for VMware ESX* virtualization testing
- Network traffic traveled across a 10 Gbps Ethernet backplane

An HP fibre channel-based storage solution hosted the data.

Single-blade Performance Benchmark Configuration

In addition to the 16-blade enterprise server configuration above, a single, eight-core server (with Hyper-Threading enabled) and without the ECP layer was tested to benchmark performance for smaller installations.

Virtualization Performance Benchmark Configuration

Built on VMware vSphere, the virtualized server configuration used a single eight-core blade to support three virtual machines (VMs), each with four virtual processors and 12 GB of memory. A single, simulated web-based client provided the user tasks. Figure 2 (shown on the next page) illustrates the virtual environment configuration.

In addition to application performance, live migration of a VM from one blade to another was tested using VMware VMotion*.

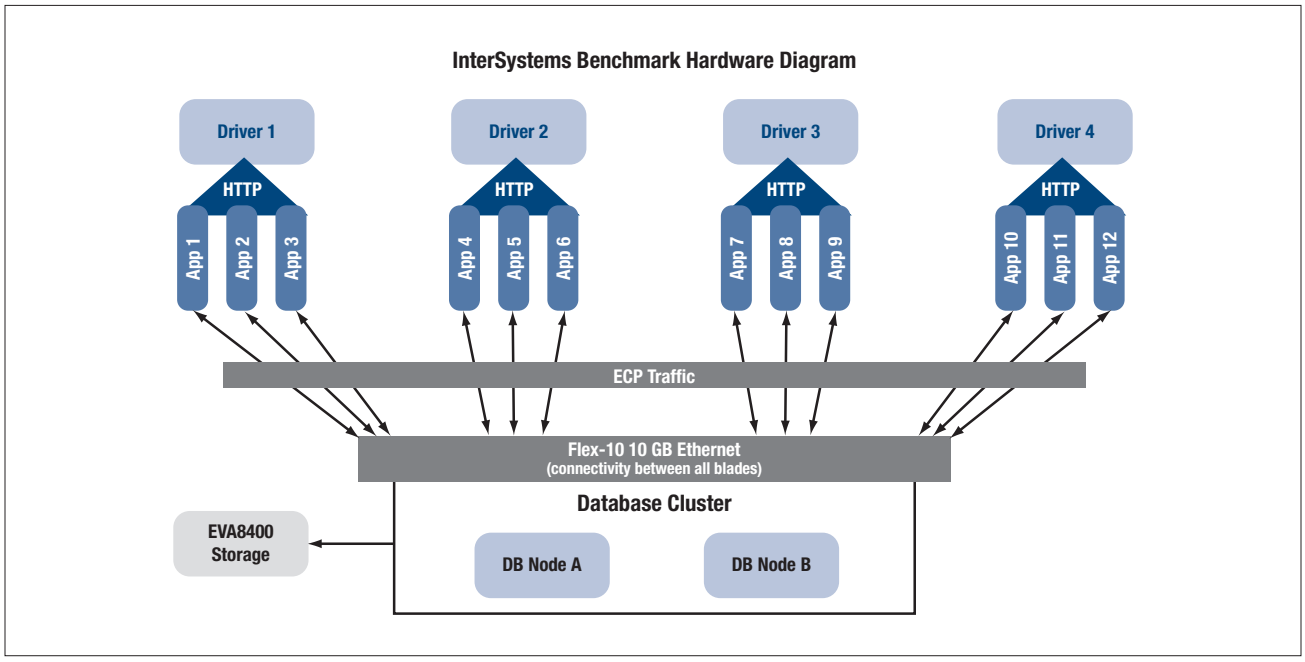

Figure 1. Scalability Benchmark Configuration

Results

The benchmark showed the configuration achieved unmatched performance and scalability for all three configurations running InterSystems Caché: a single-node, multi-core system; a multi-node, multi-core system; and a virtualized environment.

Results Summary

At 12 servers and four benchmark drivers, the system successfully sustained a peak average of 768 episodes/minute and over 11 million database accesses/second (known as database references) with maximums over 19 million database accesses/second. This activity equates to 67,200 simulated concurrent users – the population of a medium-sized city – simultaneously accessing the system adding, updating, and retrieving patient information.

Figure 3 graphs the throughput the application and system achieved during the 12 application server/768 episodes per minute test run.

Throughout the testing, component transactions achieved sub-second response times. The system not only scaled to massive proportions, it did so while maintaining optimal user response times; *user experiences were not impacted*.

Single-server Results

The single-server, native configuration, without ECP, sustained a peak average of 96 episodes/minute and 1.4 million database accesses/second. It achieved maximums near 3 million database accesses/second. This is equivalent to 8,400 concurrent users, supported by a single-blade, eight-core server with Hyper-Threading enabled. Figure 4 graphs the results of the 96 episodes/minute tests.

While InterSystems does not recommend an installation of this configuration for so many users, the benchmark illustrates that the hardware and software can certainly achieve this level of performance. Such an implementation offers high performance for supporting EMR activities in medium installations of approximately 1,000 users without impacting user response times and user productivity. Table 2 (shown on the next page) lists the single, 8-core server results.

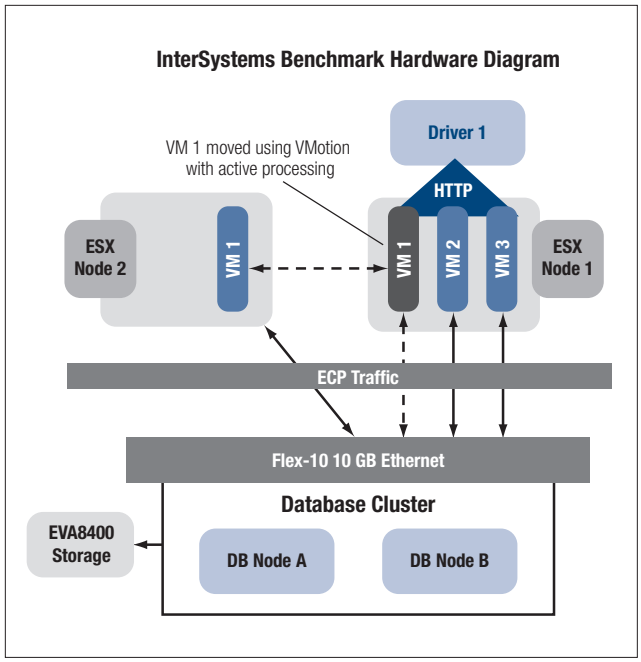

Figure 2. Virtual Environment Configuration

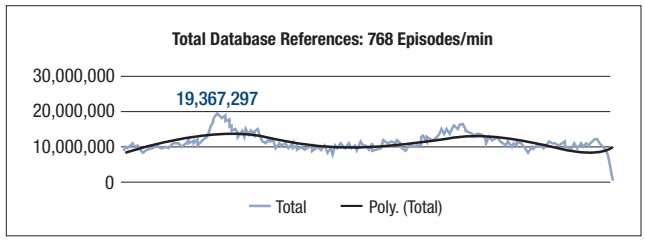

Figure 3. Databases access/second during 768 episodes/minute test run

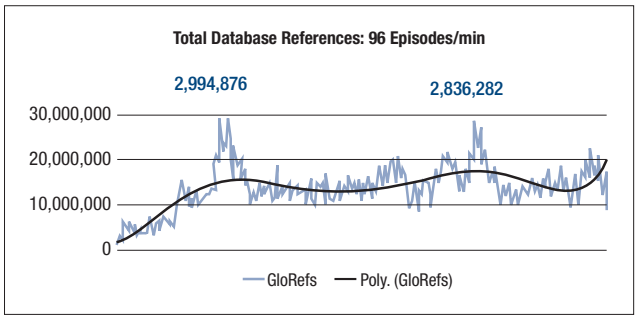

Figure 4. Databases access/second during 96 episodes/minute test run (single-blade server)

Table 2. Single Caché® Implementation Results

Configuration	Avg. Episodes/Min.	Avg. Throughput (db accesses/sec.)	Peak Throughput (db accesses/sec.)
8-core Database Server Without ECP (Hyper-Threading Enabled)	96	1,400,000	2,994,876

Table 3. Scalability Results – 1 to 12 Application Servers

Number of 8-core Application Servers	Avg. Throughput (db accesses/sec.)	Peak Throughput (db accesses/second)
1	853,769	1,363,796
4	3,532,926	4,988,669
8	7,710,205	13,690,211
12	11,100,000	19,367,297

Scalability Results

The benchmark revealed Caché linearly scales from one to 12 application servers and can support 67,200 simulated concurrent users executing typical database activities. Table 3 lists the results of scalability testing with ECP. Figure 5 charts the scalability data.

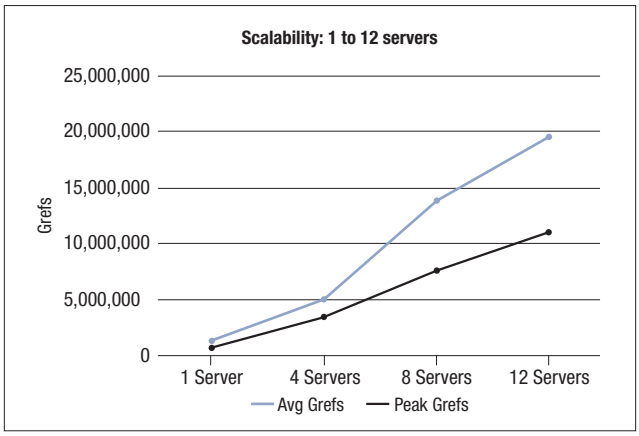

Figure 5. Scalability results chart

Network Performance

The 10 Gbps network backplane in the HP BladeSystem* proved to be more than adequate for database traffic for scalability testing. InterSystems Caché deployments typically recommend 1 Gbps Ethernet per four application servers. This indicates there is additional network capacity to further scale out the server and support a significantly larger number of users.

Virtual Environment Performance Results

With four virtual processors in each of three VMs, each virtualized system could sustain 16 episodes/minute per VM – a total of 48 episodes/min for the system. This was achieved on a single, 8-core/16-thread blade based on Intel Xeon processor X5570.

Application response times were slightly impacted, but still considered good, compared to the non-virtual environment. CPU utilization noticeably increased to support the virtual environment, but engineers believed they could scale out more VMs to support additional users if necessary. As with non-virtual benchmarks, this test revealed Caché could perform exceptionally well in a virtual environment with the proper hardware configuration. A major international installation uses this virtualized configuration today with excellent results.

Live Migration Results

Maintaining a live system, even in the presence of system component failures, is essential for business-critical applications, like EMR. Traditionally, failover solutions can be expensive with entire hot server backups ready to go online at a moments notice. Partition migration testing showed that with VMware HA VMotion, application servers could migrate seamlessly from one physical server (blade) to another without application timeouts, or impacting user response times or perception.

Encryption Burden Results

Database encryption can create additional overhead for the CPU. During single-server testing, performance was measured for both encrypted and unencrypted transactions, resulting in only approximately one percent impact on performance with encryption enabled. Keep in mind, the amount of overhead will be application-specific by the amount of physical IO occurring.

Conclusion

With a 16-blade, 2X quad-core Intel Xeon processor X5570-based blade server platform for this benchmark, InterSystems Caché database achieved a breakthrough level of performance never previously seen for a multi-core, multi-node or multi-core, single-node database solution.

Able to deliver 768 episodes/min and support over 67,000 concurrent users on a commercial, 16-blade deployment, Caché has shown scalability that can grow to support large community – and even national – EMR applications. Even a single-blade, 8-core configuration delivered unmatched performance at 96 episodes/min.

Benchmarking additionally revealed Caché supports exceptional levels of performance for virtual environments based on VMware vSphere, and can migrate VMs seamlessly across blades.

Compared to large, traditional deployments, Intel Xeon processor-based blade servers are scalable, deliver high-performance, and offer cost-effective solutions for EMR applications based on Caché databases. Caché running on the Intel Xeon processor X5570 delivers an optimized balance of scalability, cost of deployment, performance, and reliability.

Solution provided by:

[1] Episodes per minute (episodes/minute) is the metric used by InterSystems to gauge throughput performance of the application being tested. An episode is equivalent to a patient admission or visit through the EMR system. Extensive performance analysis of live EMR sites shows episodes/minute directly correlates to CPU utilization. InterSystems engineers use a baseline of 16 episodes/minute, which was recorded at a live site during peak usage of this large installation. Using this observation, episodes/minute relates to concurrent users by using an existing customers specific transaction or component mixture during peak usage periods and scaled it accordingly. Database references/second is also reported for comparison purposes.

[2] Intel® Virtualization Technology requires a computer system with an enabled Intel® processor, BIOS, virtual machine monitor (VMM) and, for some uses, certain platform software enabled for it. Functionality, performance or other benefits will vary depending on hardware and software configurations and may require a BIOS update. Software applications may not be compatible with all operating systems. Please check with your application vendor.

0810/SG/OCG/XX/PDF

Please Recycle

Order Number: 324075-001US