

# Data Scalability with InterSystems Caché® and Intel® Processors

InterSystems Caché® 2015.1 soars from 6 million to more than 21 million end-user database accesses per second on the Intel® Xeon® processor E7 v2 family compared to Caché 2013.1 on the Intel® Xeon® processor E5 family



“The scalability and performance improvements of Caché® version 2015.1 are terrific. Almost doubling the scalability, this version provides a key strategic advantage for our user organizations that are pursuing large-scale medical informatics programs as well as aggressive growth strategies in preparation for the volume-to-value transformation in healthcare.”

– Carl Dvorak,  
President, Epic

## Overview

With data volumes soaring and the opportunities to derive value from data rising, database scalability has become a crucial challenge for a wide range of industries. In healthcare, the rising demands for healthcare services and significant changes in the regulatory and business climates can make the challenges particularly acute. How can organizations scale their databases in an efficient and cost-effective way?

The InterSystems Caché® 2015.1 data platform offers a solution. Identified as a Leader in the Gartner Magic Quadrant for Operational Database Management Systems,<sup>1</sup> Caché combines advanced data management, integration, and analytics. Caché 2015.1 is optimized to take advantage of modern multi-core architectures and represents a new generation of ultra-high-performance database technology. Running on the Intel® Xeon® processor E7 v2 family, Caché 2015.1 provides a robust, affordable solution for scalable, data-intensive computing.

To examine the scalability of Caché 2015.1, InterSystems worked with performance engineers at Epic, whose electronic medical records (EMRs) and other healthcare applications are deployed by some of the world’s largest hospitals, delivery systems, and other healthcare organizations. The test team found that Caché 2015.1 with Enterprise Cache Protocol® (ECP®) technology on the Intel Xeon processor E7 v2 family achieved more than 21 million end-user database accesses per second (known in the Caché environment as Global References per Second or GREFs) while maintaining excellent response times. This was more than triple the load levels of 6 million GREFs achieved by Caché 2013.1 on the Intel® Xeon® processor E5 family.

## Caché 2015.1: Focusing on Scalability

InterSystems Caché is a leading SQL/NoSQL data platform designed to deliver high performance and massive scalability. To create the next generation of breakthrough applications, solution developers require a data platform that provides multiple modes of access to operational data. Using Caché, data is simultaneously available using NoSQL, SQL, or object paradigms. Caché’s just-in-time analytics capabilities, on both

structured and unstructured data, give developers the means to offer insight at the point of action.

Caché provides frameworks for rapidly developing mobile and Web applications. Simplified operations are a hallmark of Caché-based solutions, whether deployed on-premises or in cloud or hybrid configurations. Tens of thousands of Caché applications are available in healthcare, financial services, retail/logistics, and government solutions.

Caché supports both horizontal and vertical scalability. Its ECP technology is optimized to maintain responsiveness as organizations add application servers. Caché also supports the added capability of scale-up configurations, which grow as more processors are added to a system.

With Caché 2015.1, InterSystems focused on using modern processor architectures and capabilities to increase scalability (that is, throughput) without sacrificing latency (that is, response times). Among the enhancements:

- Caché 2015.1 implements new algorithms to improve the parallelization of specific workloads. This is evident when observing application response times that remain unchanged even while adding workloads and users to the system in a scale-out model.
- Large-core-count systems often introduce non-uniform memory access (NUMA)-related slowdowns, but hardware vendors often provide chip-specific instructions and programming techniques that can mitigate the NUMA penalty. Caché 2015.1 implements similar techniques so application response times do not degrade as the system scales up.
- Several modifications optimize critical sections and internal resources to ensure much higher throughput through parallelization between processes in both scale-out and scale-up models.

### Practical Guidance Based on Real-World Data and Scenarios

To help ensure reliable, productive experiences for its users, Epic conducts rigorous tests as a basis for developing comprehensive sizing guidelines. Its test program is designed to reflect realistic, meaningful work scenarios of large healthcare enterprises. This includes Epic working with a customer

site to obtain a copy of its database to increase the validity of the testing. Potential personal health information is scrambled, and relevant privacy protections are in place during the testing.

For the testing described in this paper, Epic collaborated with Sanford Health to obtain a copy of its multi-terabyte database. Sanford Health is an integrated health system and the largest rural, not-for-profit healthcare system in the United States.

Based in South Dakota, it has locations in nine states, with 43 hospitals, 140 clinic locations, more than 26,000 employees, and 1,400 physicians. Sanford Health is also developing international clinics in Ghana, Mexico, and China. Sanford Health's rich database includes data for the full range of enterprise Epic applications.

### Methodology

Epic's and InterSystems's performance engineers performed a series of tests to validate the level of throughput (in database accesses per second or GREFs) that Caché 2015.1 could achieve while maintaining acceptable responsiveness. Test engineers also compared the scalability of Caché 2015.1 and Caché 2013.1. Tests were run in November and December 2014 at Epic's headquarters in Verona, Wisconsin. Tests used a prerelease version of Caché 2015.1 provided by InterSystems.

The testing used real Epic applications and their associated database code. Client-side user loads were synthetically created by Epic simulation tools to run against the database. These user loads:

- Simulate the second-to-second activities of EMR users who serve in various roles and carry out diverse workflows

- Scale up activity in the same way organizations do, by adding more users and application functionality
- Exercise the same code base as Epic users, with the same suite and scope of enterprise application code and the same network protocols

Epic and InterSystems engineers ran tests that stressed the database by incrementally increasing the volume of connected users—doctors, nurses, pharmacists, and others—and thus the volume of database accesses per second. As the loads increased, the test team measured system metrics and application response times at each load level.

The test team was most interested in two key metrics:

- **Database accesses per second** or GREFs are a measure of the throughput of the Caché database. Increasing the number of simulated users also increases the volume of concurrent database interactions and thus the database accesses per second.
- **Epic response time** is a measure of how quickly the database can complete the series of rich, complex database queries performed as part of the workflow step for each simulated end user.

Load levels were considered valid only if the application response times were deemed fast and system metrics remained within a specified range.

To fully assess responsiveness, live test engineers did more than simply track the Epic response times based on the client-side loads. Engineers also manually performed workflows in the applications while the simulated load was being applied. Engineers calibrated the values for the response-time metric in relation to their real-world experience.

### System Configurations

Test engineers compared scalability and response times for two configurations:

- Caché 2013.1 with a virtualized two-socket server with the Intel Xeon processor E5-2680 (2.7 GHz) as the database server
- Caché 2015.1 with a virtualized four-socket server with the Intel Xeon processor E7-4890 v2 (2.8 GHz) as the database server

Both platforms ran Red Hat Enterprise Linux\* and used VMware vSphere\* virtualization software.

Both platforms were configured to Epic's best practices for Caché ECP configurations. Both test environments had spare processor capacity available on the database and ECP tiers, ensuring this was not a bottleneck. The ECP application tier was a mixed environment with hundreds of cores of processing power. All systems were configured to ensure that hardware capacity (for example, I/O latency, bandwidth, and memory capacity) would not artificially limit scalability or slow the systems' responsiveness.

Network and storage were also over-provisioned for all tests. The storage unit supporting the test configurations was a large-scale Flash\* or solid-state drive system that was more than adequate for the task.

### Results

Figure 1 depicts the test results. It shows that Caché 2015.1 on the virtualized platform based on the Intel Xeon processor E7-4890 v2 achieved over three times more end-user database accesses per second or GREFs than Caché 2013.1 on the virtualized platform based on the Intel Xeon processor E5-2680—more than 21 million GREFs compared to approximately 6 million.

The results also show that the platform maintained excellent response times while doing so. A value of 2 or less is deemed an acceptable test with excellent response times. Note that these response times are normalized values, not given in seconds or other units but showing a comparison of how response times changed from test to test. They thus represent normalized values summarized across the simulated workflows performed at each load level.

“InterSystems Caché® has historically been a super-scalable, high-performance database. Our goal for Caché version 2015.1 was ambitious: double the throughput [scalability] without sacrificing latency [responsiveness]. Our work with Caché 2015.1 has more than paid off—not only can we boast over double the scalability, but also double the performance.”

– Mark Bolinsky,  
Technology Architect, Manager,  
InterSystems Corporation

Threefold Improvement in Database Scalability with InterSystems Caché® Enterprise Cache Protocol (ECP) and Epic® EMR Software on Intel® Xeon® Processors




Figure 1. Threefold Improvement in Database Scalability with InterSystems Caché® Enterprise Cache Protocol® (ECP®) and Epic® EMR Software on Intel® Xeon® Processors

With ECP running on the Intel Xeon processor E7-4890 v2, the test team calculated that Caché 2015.1 sustained nearly 22 million database accesses per second on a diverse application server compute environment consisting of 140 CPU cores and a total of around 440 GHz of processing power. This works out to a very impressive figure of approximately 50,000 database accesses per second per gigahertz of compute. The memory configuration was roughly 500 GB across the entire compute layer. Testing stopped at 22 million database accesses per second because the test team ran out of hardware, not because that figure represents any type of scalability wall.

Various additional factors were not accounted for in these tests but may determine how to size a system in real-world situations. Factors that influence system sizing include variance across different types of workload patterns, maintaining real-time replicated environments, and other conditions. To help its customers achieve optimal performance, Epic provides detailed sizing advice that encompasses these and other factors, as well as tests such as those described in this paper.

### Implications

InterSystems Caché 2015.1 and the Intel Xeon processor E7 v2 family offer an outstanding and affordable solution for mission-critical, data-intensive computing in healthcare and in other industries.

The Intel Xeon processor E7 v2 family combines performance and cost-effectiveness with enterprise-grade scalability. Based on Intel's most robust and powerful server technologies, the Intel Xeon processor E7 v2 family is available with up to 15 cores and 30 threads per socket. It provides up to 25 percent more cache and three times the memory capacity of previous members of the Intel Xeon processor E7 family. Intel's product road map highlights ongoing advances in the performance and capabilities of the Intel Xeon processor E7 family.

For healthcare organizations, the combination of InterSystems Caché 2015.1, Epic EMR software, and the Intel Xeon processor E7 v2 family provides an enterprise-class solution for rising volumes of EMR data, genomics data, and other sources of data growth. These three robust

technologies enable organizations to accommodate the mergers, acquisitions, and collaborations that are further swelling the size of their healthcare databases. In doing so, these technologies can position healthcare enterprises to not only deliver outstanding care, but also to make the most of the changes in the regulatory and business climate.

Data is a foundation of competitive business advantage. See how InterSystems and Intel can help you make the most of your data—to improve healthcare with Epic EMR applications or to transform information into insights in any industry.

### Learn More

**InterSystems Caché Data Platform:**  
[www.intersystems.com](http://www.intersystems.com)

**Epic:** [www.epic.com](http://www.epic.com)

**Intel Xeon Processor E7 Family:**  
[www.intel.com/content/www/us/en/processors/xeon/xeon-processor-e7-family.html](http://www.intel.com/content/www/us/en/processors/xeon/xeon-processor-e7-family.html)

**Sanford Health:** [www.sanfordhealth.org](http://www.sanfordhealth.org)

**Twitter:** @InterSystems, @IntelHealth, @IntellITCenter



<sup>1</sup> InterSystems Recognized As a Leader in Gartner Magic Quadrant for Operational DBMS, October 16, 2014. <http://www.intersystems.com/our-products/cache/intersystems-recognized-leader-gartner-magic-quadrant-operational-dbms/>

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. **For more information go to** [www.intel.com/performance](http://www.intel.com/performance)

Intel does not control or audit the design or implementation of third party benchmark data or Web sites referenced in this document. Intel encourages all of its customers to visit the referenced Web sites or others where similar performance benchmark data are reported and confirm whether the referenced benchmark data are accurate and reflect performance of systems available for purchase.

This document and the information given are for the convenience of Intel's customer base and are provided "AS IS" WITH NO WARRANTIES WHATSOEVER, EXPRESS OR IMPLIED, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT OF INTELLECTUAL PROPERTY RIGHTS. Receipt or possession of this document does not grant any license to any of the intellectual property described, displayed, or contained herein. Intel® products are not intended for use in medical, lifesaving, life-sustaining, critical control, or safety systems, or in nuclear facility applications.

Copyright © 2015 Intel Corporation, Epic Systems Corporation, InterSystems. All rights reserved. Intel, the Intel logo, Intel Inside, and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries.

InterSystems, Caché, and ECP are trademarks or registered trademarks of InterSystems Corporation.

Epic is a registered trademark of Epic Systems Corporation.

