WHITE PAPER
Intel® Advanced Encryption Standard New Instructions
Enterprise Security

High-Performance Encryption for Databases in Financial Services

Using Intel® Advanced Encryption Standard New Instructions with InterSystems Caché* Substantially Improves Encryption Performance and Reduces Computational Overhead

EXECUTIVE SUMMARY

Financial services companies have an ever-growing need to encrypt databases containing sensitive customer and trade data. However, using encryption on these databases can require significant computational resources, potentially impacting trading latencies. Intel® Advanced Encryption Standard New Instructions (Intel® AES-NI), included in the Intel® Xeon® processor 5600 and E5 product families (and more recent Intel Xeon processor families), accelerates encryption and greatly reduces computational overhead.

> "Encryption is now a requirement for data in many domains. We've had an efficient implementation, and now, with the hardware assist from Intel, we can offer encryption with much lower impact on overall system performance."
>
> *–Robert Nagle*
> *Vice President of Software Development, InterSystems Corporation*

Key Findings

Tests show that even as data and computational time increase, there is negligible degradation in performance for encryption or decryption. Use of Intel AES-NI with Caché's interleaved cipher blocks can speed up encryption by a factor of 20 or more.

Securing Sensitive Data to Prevent Breach

A global security study by Deloitte Touche Tohmatsu Limited found that in 2011, one in four financial services firms suffered security breaches.[1]

The cost to businesses of exposing data such as Social Security and credit-card numbers climbed 7 percent in 2011 to an average of USD 7.2 million per incident according to a report by Ponemon Institute LLC, an information security research group.[2]

To mitigate this type of risk, organizations are increasingly using encryption on their databases. However, encryption can require additional computational resources. This can impact application scalability and increase the cost of deployment.

The right combination of database platform, high-end processors, and encryption/decryption solution can enable cost-effective, scalable security for high-performance, mission-critical systems.

Intel AES-NI for Encryption and Decryption

Intel AES-NI implements strong encryption and decryption while greatly reducing the processing time required.[3] It consists of a new set of seven instructions, four of which implement the core of the AES algorithm and accelerate data encryption and decryption on the Intel Xeon processor 5600 and E3, E5, and E7 product families and more recent Intel Xeon processor families.

By accelerating performance, Intel AES-NI provides improved application scalability and more affordable data protection. These benefits are maximized by using Intel AES-NI in a mode that interleaves the processing of multiple cipher blocks.[4]

InterSystems Caché for Financial Services

InterSystems Caché is a high-performance database and rapid application development and deployment platform that powers many large-scale systems in financial services. The Caché high-performance database engine stores data in highly compressed, sparse, multi-dimensional arrays called "globals." Application programmers can access and manipulate this data through objects, SQL*, or direct access to the underlying structure.

Caché database files consist of sequential fixed-size blocks of 8, 16, 32, or 64 kilobytes. Blocks contain application data, indices, or metadata such as directories, pointers, or allocation maps. A single Caché instance can have many databases, with each database using one of these block sizes and each stored in a separate disk file. When in use by applications, database block images are stored in shared memory buffers.

Caché distributes application processing among many processes. When a process needs data that is in a database block not currently resident in a shared memory buffer, that process allocates a free buffer and reads the needed database block from disk. Other processes are not affected.

When the numbers of modified, unmodified, and free shared memory buffers hit any of several thresholds, or a maximum time interval elapses, a set of write daemon processes writes the contents of the modified buffers back to disk. This allows for efficient bulk processing of disk writes and allows application processes to continue while disk writes are in progress.

Reducing Computational Overhead

Databases that contain sensitive data should be encrypted to prevent compromise. The Caché database encryption feature was designed to:

- **Encrypt the entire contents** of the database, including all structural metadata, except for the single initial label block.
- **Perform size-preserving encryption** that maintains the efficient mapping of database blocks to disk hardware.
- **Ensure the encryption and decryption operations** are completely transparent to applications and the database engine.
- **Ensure identical data** stored in different database blocks is encrypted differently, minimizing information leakage.
- **Ensure encrypted databases** are completely portable between Caché instances running on different hardware and OS platforms.

Caché database encryption uses the AES algorithm in Cipher Block Chaining (CBC) mode at the database block level, with an initialization vector derived by encrypting the database block number.[5,6] Database encryption and decryption are performed at the interface between the Caché engine and the operating system file system.

Block-level database encryption provides high performance by amortizing fixed initialization overhead over large blocks. The choice to encrypt fixed-sized database blocks also allows for optimal exploitation of instruction-level hardware encryption. The AES algorithm operates on 16-byte cipher blocks, performing 10 to 14 rounds (depending on the key length) of multiple cryptographic primitives on each cipher block, using a different 128-bit round key for each round. Intel AES-NI instructions perform one complete round on one cipher block as a single instruction.

Because the Intel Xeon processor implements a pipelined architecture and Intel AES-NI instructions take multiple clock cycles to complete, multiple cipher blocks must be interleaved at each round for maximum processor utilization. In CBC mode, the processing of each cipher block requires the encrypted data from the previous cipher block. Cipher block interleaving is, therefore, straightforward for decryption, when all the encrypted cipher blocks are initially available and multiple cipher blocks can be interleaved in the instruction pipeline. It is not possible for encryption in the general case. Because Caché encrypts the contents of a large pool of fixed-sized shared memory buffers during write daemon processing, it can interleave cipher blocks from multiple buffers to achieve the same CPU utilization for encryption as for decryption. Caché automatically detects the Intel Xeon processor product family on which it is running and optimizes the interleaving factor to maximize performance.

Case Study: Encryption Performance Testing

Intel and InterSystems measured the computational overhead of Caché AES-CBC database encryption and decryption. Measurements were made on an unencrypted database, an encrypted database using an optimized AES-CBC software implementation (software), and an encrypted database using an interleaved AES-CBC implementation with Intel AES-NI hardware (interleaved Intel® AES-NI).

Figure 1 shows the results of these measurements on a 2.93 GHz Intel Xeon processor X5670. Figure 2 shows the results on a 2.7 GHz Intel Xeon processor E5-2680. The benefits of hardware encryption are striking. While unencrypted data will always process faster, with Intel AES-NI, there is little degradation in performance, even as time and data increase. Table 1 lists the computational overhead and speed-up factors by test case and processor.

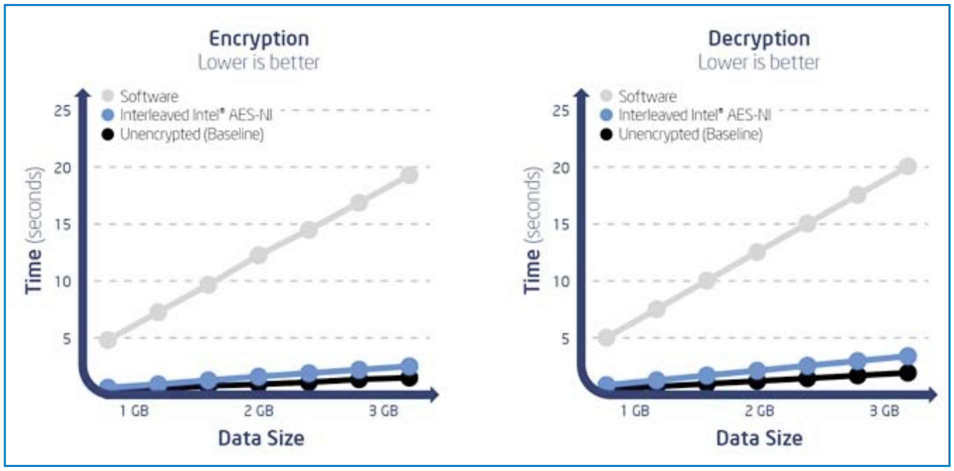

Figure 1. Performance Results of Encryption (Left) and Decryption (Right) on a 2.93 GHz Intel® Xeon® Processor X5670

Note: "Software" refers to the optimized AES-CBC software implementation. "Interleaved Intel AES-NI" is the interleaved AES-CBC implementation with Intel AES-NI hardware.

In addition to basic tests, which were focused solely on encryption and decryption speeds, testing was done on a simple but common financial services application.

The application consisted of inserting a substantial number of client accounts and their associated positions, and then iterating though them. This is a common task used in risk analysis for calculating net positions and other purposes. The application was run with a variable number of clients (1 to 5 million) with 20 positions for each client. Inserts were handled by a Java* program running 10 threads, each handling one-tenth of the load. Ten Java threads were created, each of which iterated through the entire record set. All work was done on a single computer.

The full database of 5 million clients and 100 million positions was written to an encrypted database in just over 37 minutes, at a rate of over 2,200 clients per second and over 44K positions per second. This performance was better than 99 percent compared to writing to the unencrypted database. The primary performance factor on the system was the speed of the disk drive.

Queries were run against different sized data sets. Database memory cache was set to a low 500 MB (compared to over 29 GB of data), which required Caché to continually go to disk to resolve queries, forcing a large amount of decryption of data. Despite this, query performance with the Intel AES-NI hardware was better than 96 percent of performance compared to running against an unencrypted database.

This should be considered a worst-case scenario. In most applications—with larger database caches and more re-querying of the specific records—the performance difference between unencrypted and encrypted would be even less. In fact, this scenario was picked because it would stress Caché encryption abilities more than just about any other common use case in financial services.

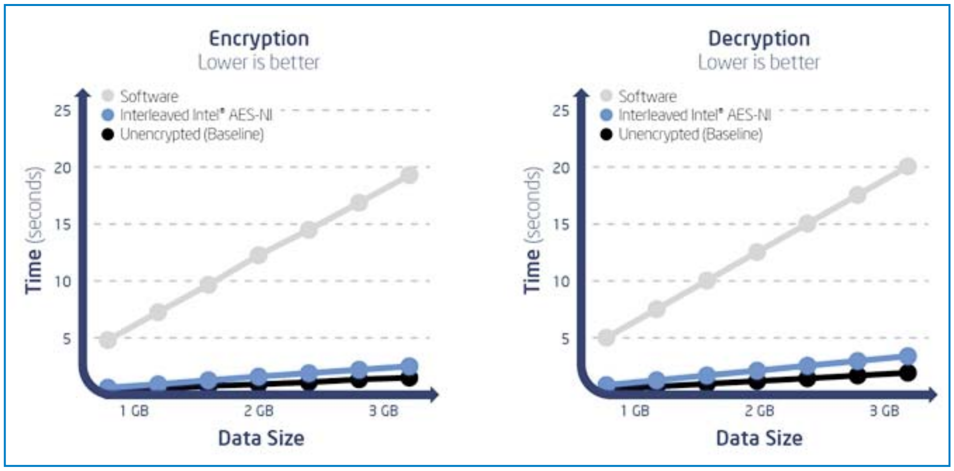

Figure 2. Performance Results of Encryption (Left) and Decryption (Right) on a 2.7 GHz Intel® Xeon® Processor E5-2680

Note: "Software" refers to the optimized AES-CBC software implementation and "interleaved Intel® AES-NI," the interleaved AES-CBC implementation with Intel AES-NI hardware.

Table 1. Observed Computational Overhead and Speed-Up Factors Translated into Clock-Speed Invariant Units

Implementation		Intel® Xeon® Processor X5670		Intel® Xeon® Processor E5-2680	
		Computational Overhead (Clocks/Byte)	Speed-Up Factor	Computational Overhead (Clocks/Byte)	Speed-Up Factor
Encryption	Software	16.0		12.0	
	Interleaved Intel AES-NI	1.0	16x	0.5	24x
Decryption	Software	16.0		12.0	
	Interleaved Intel AES-NI	1.3	12x	0.8	15x

Performance test results show Intel® AES-NI tracks nearly identical to the baseline of unencrypted data.

Conclusion

The design of InterSystems Caché database encryption allows maximally efficient utilization of the Intel AES-NI capability implemented in the latest Intel Xeon processor families. Intel AES-NI allows Caché to perform encryption and decryption faster, enabling it to handle an ever-increasing amount of sensitive data at reduced hardware cost. In addition to servers and backups, Intel AES-NI hardware-assisted encryption and decryption can also be useful for client platforms and encryption of data in transit.

For more on InterSystems Caché, see the **InterSystems Caché Technology Guide**.

Find more on Intel Advanced Encryption Standard New Instructions **here.**

Factors to Consider For Database Encryption Solutions

There are several factors to consider make use of this advanced capability.

Be aware of regulations, data protection laws, and breach notification rules for your business' geographic location(s) and requirements for protecting confidentiality of sensitive financial data.

Be aware of the increasing risk of breaches from compromised servers and address these types of risks in your risk assessments.

Consider server database encryption a key safeguard as part of a holistic approach to protect confidentiality of sensitive financial data, and avoid breach.

If database encryption is part of your solution, consider the advantages of deploying InterSystems Caché on hardware that supports Intel AES-NI.

[1] "2012 Global Financial Services Industry Security Study: Breaking Barriers," http://www.deloitte.com/security-survey-2012.

[2] "2010 U.S. Cost of a Data Breach," Ponemon Institute (March 2011).,
http://www.symantec.com/about/news/resources/press_kits/detail.jsp?pkid=ponemon.

[3] Intel Advanced Encryption Standard New Instructions (Intel® AES-NI) is a set of instructions that consolidates mathematical operations used in the Advanced Encryption Standard algorithm. Enabling Intel AES-NI requires a computer system with an Intel AES-NI-enabled processor as well as non-Intel software to execute the instructions in the correct sequence. For availability of Intel AES-NI enabled processors or systems, check with your reseller or system manufacturer.

[4] Secure Cloud with High Performing Intel Data Protection Technologies. www.youtube.com/watch?v=I0ALeQjS7FA].

[5] Federal Information Processing Standards Publication 197, "Specification for the Advanced Encryption Standard.,"
http://csrc.nist.gov/publications/fips/fips197/fips-197.pdf.

[6] National Institute of Standards and Technology. Special Publication 800-38A, "Recommendation for Block Cipher Modes of Operation," http://csrc.nist.gov/publications/nistpubs/800-38a/sp800-38a.pdf.

Intel, Xeon, and the Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

*Other names and brands may be claimed as the property of others.

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

A "Mission Critical Application" is any application in which failure of the Intel Product could result, directly or indirectly, in personal injury or death. SHOULD YOU PURCHASE OR USE INTEL'S PRODUCTS FOR ANY SUCH MISSION CRITICAL APPLICATION, YOU SHALL INDEMNIFY AND HOLD INTEL AND ITS SUBSIDIARIES, SUBCONTRACTORS AND AFFILIATES, AND THE DIRECTORS, OFFICERS, AND EMPLOYEES OF EACH, HARMLESS AGAINST ALL CLAIMS COSTS, DAMAGES, AND EXPENSES AND REASONABLE ATTORNEYS' FEES ARISING OUT OF, DIRECTLY OR INDIRECTLY, ANY CLAIM OF PRODUCT LIABILITY, PERSONAL INJURY, OR DEATH ARISING IN ANY WAY OUT OF SUCH MISSION CRITICAL APPLICATION, WHETHER OR NOT INTEL OR ITS SUBCONTRACTOR WAS NEGLIGENT IN THE DESIGN, MANUFACTURE, OR WARNING OF THE INTEL PRODUCT OR ANY OF ITS PARTS.

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 http://www.intel.com/performance.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined". Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request. Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

0713 ER/SS