



Solution Brief

Hyperscan Pattern Matching Software
Intel® Architecture Processors

HIGHLY-SCALABLE DPI (PATTERN MATCHING) PERFORMANCE ACROSS INTEL® PROCESSORS USING HYPERSCAN

Hyperscan optimizes content inspection performance on Intel® architecture, scaling from Intel® Atom™ processors to Intel® Core™ processors to Intel® Xeon® processors

Executive Summary

Combating the growing amount of malware is becoming an ever-increasing, resource-intensive task requiring the deployment of even more advanced scanning capabilities. Content scanning technologies are supported on a wide variety of applications and equipment types, including large cloud-based server blades, security appliances, switches, and routers. As an alternative to using custom ASICs and equivalent hardware to perform the task of pattern matching, equipment designers can now address the need with a simplified software-based approach. Hyperscan is a software pattern matching library that fully scales across Intel® architecture to deliver the highest levels of content inspection performance as demanded by today's security applications.

Pattern matching is the underlying function at the heart of most security applications. To drive performance and scaling, this technology has typically relied on purpose-built or dedicated hardware: a design approach that often leads to complex development and high product costs. In fact, the industry is rapidly moving away from costly,

dedicated compute nodes to software-driven architectures using network functions virtualization (NFV) and software-defined networking (SDN). Intel's Hyperscan pattern matching solution is ideal for NFV/SDN-based equipment, offering a highly flexible and scalable content inspection capability. Hyperscan performance and functionality, whether virtualized or non-virtualized, scales linearly on a per core/thread basis on Intel silicon.

This paper reviews the content inspection performance benchmark data and demonstrates Hyperscan's ability to deliver scalable pattern matching throughput performance when running on entry-level Intel® Atom™ processors to high-end Intel® Xeon® processors. Security vendors can use the data to better characterize pattern matching performance of various Intel architecture processor SKUs based on their operating frequency, number of processor cores, and L3 cache size.

“Pattern matching is the underlying function at the heart of most security applications.”

Deep Packet Inspection

Pattern matching is a complex technique and involves scanning large amounts of data against a database of patterns (rule sets) in order to detect and identify threats. The deeper the inspection, the greater the packet processing requirements, which ultimately impacts the performance of the security application. For example, widely used applications such as Firewalls, Intrusion Prevention Subsystems (IPS), and Unified Threat Management (UTM) have become highly resource intensive, often creating performance bottlenecks at critical points in the network. Therefore, the performance engineering of applications such as these has become a priority.

Hyperscan

HyperScan is a software pattern matching library that can match large groups of regular expressions against blocks or streams of data, ideal for applications that need to scan large amounts of data at high speed. Hyperscan provides a simple API that is easy to integrate and is a drop-in replacement for libPCRE to deliver scan

performance that is orders of magnitude better. When deployed on an Intel processor-based platform, Hyperscan takes advantage of features such as hyperthreading, receive side scaling, and SIMD instructions to provide optimized scanning performance of over half a terabit per second on high-end Intel Xeon processors. In addition, cache-rich Intel architecture allows large matching tables to remain in cache during scanning, thus keeping memory-access overhead to a minimum.

Scanning Intelligence

Hyperscan's simplest use-case is a block scanning application. Such an application scans a single contiguous block of data with a set of regular expressions and collects any matches that occur. For these cases, Hyperscan provides a block mode interface that does not store state information and returns all of the matches before it completes. Many applications operate on data that may not be available as a single block. For example, network traffic scanning applications are often unable to hold all of the packets that make up a message in memory, and simply

scanning each packet ignores matches that straddle packet boundaries. To support those cases, Hyperscan also provides a streaming API, enabling such applications to easily implement cross-packet inspection. In streaming mode, the application can pass a stream of data blocks to Hyperscan, one at a time, and Hyperscan will return matches as they occur, even matches that cross the boundaries between these blocks.

Streaming support is a first class citizen for Hyperscan; matching is supported across an arbitrary number of block writes, and the full complement of supported PCRE constructs can be used. The streaming operation requires a small fixed-size stream record to store the state associated with each stream, and Hyperscan provides an easy-to-use set of interfaces for manipulating these records.

Linear Performance Scaling

Hyperscan's multi-threaded architecture takes advantage of symmetric multithreading to scale performance linearly with the number of processor cores used. Each scan

Intel® Processor	CPU Freq (Base/Turbo GHz)	Platform Details			Peak Scan Perf (Gbps)	Single Core or Thread Scan Perf (Gbps)	Approx per-core clock-for-clock perf (Gbps): scaled to 2Ghz
		Sockets	Cores/Threads	L3 Cache (MB)			
Intel® Xeon® Processor E5-2699 v3	2.3/3.6	2	36/72 (total)	45	555	21.7	18.8
Intel® Xeon® Processor D-1540	2.0/2.6	1	8/16	12	86	11.1	11.1
Intel® Xeon® Processor E3-1285 v3	3.6/4.0	1	4/8	8	76	16.3	9.0
Intel® Atom™ Processor C2758	2.4/N.A	1	8/8	0	22	1.8	1.5

Table 1. Performance Data from Hyperscan Running on Intel® Atom™ and Intel® Xeon® Processors



runs independently of the other scans, allowing for concurrent processing of different data streams without adverse performance impact.

With its ability to recompile large pattern databases into a small memory footprint, Hyperscan also helps vendors dramatically reduce memory requirements. In fact, for smaller databases it is possible for Hyperscan to take advantage of the memory rich cache architecture provided by Intel® processors to perform the scanning in-cache. The technologies significantly reduce the amount of shared memory contention in multi-core systems.

The scalable performance of Hyperscan is demonstrated in Table 1, where the peak scan performance of Intel Atom and Intel Xeon processor-based platforms ranges from 22 to 555 Gbps. Security vendors can dial in a specific cost-performance point by choosing among Intel Xeon, Intel® Core™, and Intel Atom processors with varying CPU frequencies, numbers of cores, cache size, and socket per board. This scalability spans from entry-level customer-premises equipment to high-throughput data centers, enabling security to address multiple markets with a single pattern matching product.

The test case was a database of 250 synthetic patterns composed of a variety of regular expression constructs, intended to simulate a mix of real-world patterns. The input was taken from real HTTP traffic, captured and played back from a PCAP file. The processors were 100 percent utilized in non-streaming modes. Results for streaming modes were approximately two percent lower.

Reducing Development Costs with Scalable DPI Solution

Network security vendors are looking for agile platforms that provide predictable DPI performance, and higher levels of scalability and flexibility. This is possible with Hyperscan software running on Intel processors. An equipment vendor



can integrate Hyperscan into a system software release for a particular product line and, with one integration cycle, can scale the same feature set, functionality, and API across the entire product suite from the lowest-end product to the largest multi-Gbps network server equipment. With feature consistency and performance calibration at the per core/thread level, equipment designers can streamline their design complexity while optimizing performance on a per core count basis irrespective of the product being low or high end.

About Wind River*

Hyperscan is available through Wind River*, a wholly owned subsidiary of Intel Corporation (NASDAQ: INTC), and a world leader in delivering software for the Internet of Things. The company has been pioneering computing inside embedded devices since 1981, and its technology is found in nearly 2 billion products. Wind River offers a comprehensive portfolio of solutions for addressing the system-level challenges and opportunities of IoT that is backed

by world-class global professional services, award-winning customer support, and a broad partner ecosystem. Wind River delivers the technology and expertise that enables the innovation and deployment of safe, secure, and reliable intelligent systems. To learn more, visit Wind River at www.windriver.com.

For more information about Intel security solutions for communications and enterprise infrastructure, visit <http://www.intel.com/content/www/us/en/communications/communications-enterprise-security.html>

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. UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

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. Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting Intel's Web site at www.intel.com.

For more complete information about performance and benchmark results, visit www.intel.com/benchmarks.

Copyright © 2015 Intel Corporation. All rights reserved. Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.

* Other names and brands may be claimed as the property of others. Printed in USA 0615/SG/ICMCSW/PDF ♻️ Please Recycle 332765-001US

