



# Delivering 36 Gbps DPI (Pattern Matching) Throughput on the Intel® Atom™ Processor C2000 Product Family using Hyperscan

Hyperscan optimizes DPI performance on Intel® architecture, scaling from Intel® Xeon® processor to Intel® Core™ processor to Intel® Atom™ processor



Pattern matching is at the heart of most security applications

## 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 Intel® architecture to deliver the highest levels of pattern matching performance for a best in class DPI (pattern matching) solution.

Pattern matching is used in most security applications. To drive performance and scaling, this technology typically requires purpose-built or dedicated hardware: a design approach that often leads to high development and product costs. In fact, the industry is moving away from costly, dedicated compute nodes to software-driven architectures using network functions virtualization (NFV) and software-defined networking (SDN). Intel's solution is ideal for NFV/SDN-based equipment, offering a highly flexible and scalable content inspection solution. Hyperscan performance and functionality, whether virtualized or non-virtualized, scales linearly on a per core/thread basis on Intel silicon.

This paper reviews the performance benchmark of Hyperscan running on the Intel® Atom™ processor C2000 product family, providing up to 36 Gbps DPI (pattern matching) throughput: the use case was based on scanning real world HTTP traffic against a tier-1 IPS pattern database. This solution presents a compelling price/performance position for low-end security appliances, blades, and other small form factor designs requiring advanced security functionality.

## 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 intrusion prevention (IPS) and unified threat management (UTM) have become highly resource intensive and therefore performance engineering has become a priority.

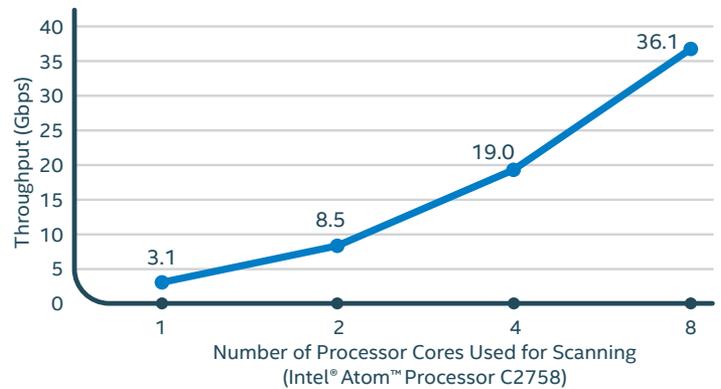
## Hyperscan

Hyperscan is an OS-independent, multithreaded software pattern matching library. With a simple API that is easy to integrate, Hyperscan is a drop-in replacement for libPCRE, but providing performance that is orders of magnitude better. When deployed on an Intel platform, Hyperscan takes advantage of features such as hyperthreading, receive side scaling, and SIMD instructions to provide optimized scanning performance of over 200 Gbps on high-end Intel® Xeon® processors.

## 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

**Figure 1.**  
Hyperscan  
Performance  
Scales Linearly  
with the Number  
of Processor Cores  
Used to Scan

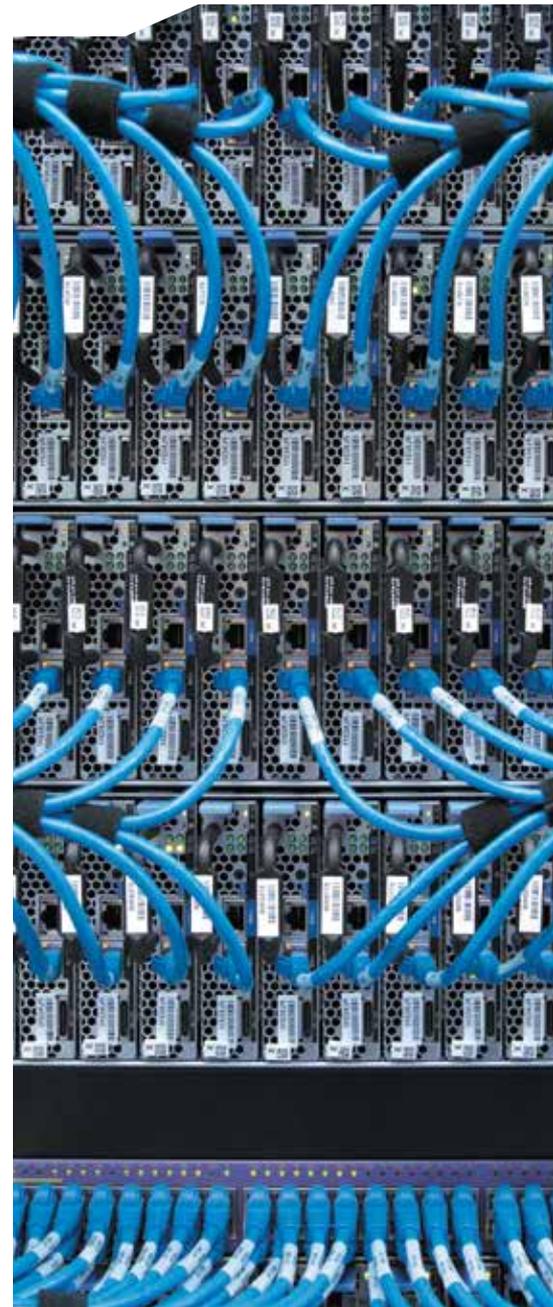


store the state associated with each stream, and Hyperscan provides an easy-to-use set of interfaces for manipulating these records.

## 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 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, leading to a more linear progression of the performance curve as the number of processor cores increase. This linear performance increase is illustrated in Figure 1, where the throughput scales from 3 Gbps to 36 Gbps as the number of cores assigned to scanning increases from one to eight on an Intel® Atom™ processor C2758.



VENDOR INTRUSION PREVENTION SOFTWARE (IPS) PATTERNS USING HTTP TRAFFIC	HYPERSCAN THROUGHPUT (Gbps) BY NUMBER OF PROCESSOR CORES			
SIGNATURE SET TYPE	1	2	4	8
Streaming, 69 Complex Signatures	3.1	8.5	18.5	36.1
Streaming, 142 Complex Signatures	2.0	5.4	12.5	25.4
Streaming, 43 Complex Signatures	0.8	1.9	4.3	10.2
Streaming, 235 Complex Signatures	0.4	0.9	1.9	4.1
Block, 13K Medium-Complexity Signatures	0.9	1.7	3.4	6.4
Block, 8K Medium-Complexity Signatures	1.1	2.1	4.2	8.0

**Table 1.** Hyperscan Performance on the Intel® Atom™ Processor C2758

### Benchmarking the Intel® Atom™ Processor C2000 Product Family

Pattern matching performance measurements can be influenced by a number of factors, including:

- The types and numbers of signatures
- The content of the incoming traffic
- The number of matches or partial matches found in the data

Therefore, benchmarking tests must use real signatures and real network traffic in order for the results to be meaningful. This was the case when Intel engineers performed benchmark testing on the Hyperscan library running on an eight-core Intel Atom processor C2758 using a complete set of current IPS signatures sourced from a leading security equipment vendor. A simple application simulated the

behavior of a real network application by reading into memory actual HTTP traffic from a PCAP file and invoking the Hyperscan APIs packet by packet.

Data was matched in streaming mode for cases where the threats might span multiple packets, and in non-streaming mode for threats that could be contained within a single chunk of data. The benchmarking application specifically measured the raw pattern matching performance, excluding the time spent in reading the PCAP file and in pre- and post-scan processing. For this benchmark, all the data used for pattern matching was resident in memory.

The benchmark results in Table 1 show near linear scalability up to eight cores for various signature set types, with raw DPI scan performance reaching up to 36 Gbps.

### 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 pattern matching 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, utilize the same DPI technology across the entire product suite from the lowest-end product to the largest multi-Gbps network equipment. With feature consistency and performance calibration at the per core 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.

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](http://www.intel.com).

Copyright © 2014 Intel Corporation. All rights reserved. Intel, the Intel logo, Intel Atom, Intel Core, and Xeon 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 0814/LL/CS/SD/PDF ♻️ Please Recycle 330943-001US

