

Scalable Security Solutions
Check Point® Open Performance Architecture
Quad-Core Intel® Xeon® Processors

Delivering Application-Layer Security at Data Center Performance Levels

Security software designed for multi-core platforms to protect infrastructure investments

Companies have always faced a tradeoff with network security. Do they lock down the network and face performance issues, or do they focus on a high level of performance at the expense of preventing possible attacks? Today, these decisions are even harder. Application-layer threats are increasingly the vector of choice for hackers and malware. These attacks—disguised as legitimate traffic—require a deeper and more sophisticated level of inspection that demands more processing power. At the same time, companies are beginning to transition to 10 gigabit (Gb) Ethernet networks, which require security solutions capable of sustaining these speeds in order to maintain a positive return on investment.

The Check Point® Open Performance Architecture security software running on Intel® multi-core processors was designed with these two challenges in mind. Rather than make companies choose between performance and security, this solution increases throughput while simultaneously raising security levels. Check Point security software provides the ability to deliver both security and performance on the same platform by providing several layers of acceleration technologies. These layers work together with advanced technologies, such as quad-core Intel® Xeon® processors, to deliver on the promise of application-layer security at high performance levels. This white paper explains the philosophy behind the Check Point Open Performance Architecture and reviews the Intel technologies used to secure service provider, enterprise and small and medium-sized business networks.



www.intel.com



Check Point®
SOFTWARE TECHNOLOGIES LTD.

www.checkpoint.com

A Delicate Balance: Performance and Security

Securing a network is a constant tradeoff between enabling users to readily access data while protecting them from cybercriminals. Users want and expect instant access to data, to systems and to other people, as in the case of Voice over Internet Protocol (VoIP). On the other hand, security is about limiting unfettered access, thereby securing data and systems and keeping them virus-free. New regulations and a higher level of security awareness have forced organizations to reexamine their security policies and place more emphasis on security. As organizations become increasingly strict about security, the degree of user access normally decreases.

Complicating matters is the fact that as security controls are increased, the security tools themselves come under a higher workload, which reduces performance and indirectly affects access levels. To protect against today's risks of highly advanced attacks and information leakage, it's necessary to perform a higher level of inspection on traffic passing through the perimeter gateway. When more security checks are placed on information, the security tools themselves face a greater processing load to implement the security policy—effectively slowing down security inspection.

This problem of balancing information access and security is evident in two key areas: increased bandwidth requirements and rising levels of application-layer threats.

Increased Bandwidth Requirements

Networks are transitioning from 1 Gb to 10 Gb Ethernet. Although this will not immediately translate into increased throughput requirements at the perimeter, security performance requirements will increase on the whole. Besides existing at the perimeter, integrated firewall/VPNs play an important role in separating network segments and segregating important servers in large offices and data centers. Internally, the change to 10 Gb will require firewalls to scale appropriately as well.

Increased Application-Layer Threats

Today, many attacks are masquerading as legitimate application-layer traffic, enabling them to threaten a whole host of applications: instant messaging, chat, peer-to-peer and Web applications, just to name a few. The reasoning behind these attacks is that traditional firewall-based security focuses on network-layer access, preventing people from accessing specific IP addresses or networks unless authorized. Modern attacks mean that a supposedly trusted user is disguising the traffic so that it passes the firewall. From a security inspection viewpoint, the answer is to perform a deeper

level of inspection, similar to intrusion prevention, on the firewall to detect application-layer threats. However, every additional security screening that is done decreases the ability for the firewall to efficiently process the traffic, slowing down its predictable performance.

The traditional method of dealing with increased security performance requirements has been to develop a closed architecture based on application-specific integrated circuits (ASICs) or other specialized hardware. These purpose-built devices are designed to efficiently handle specific tasks much faster than general-purpose processors. For some security tasks, such as network address translation (NAT) or basic packet filtering, these closed systems provide a simple way to accelerate security performance.

The problem with closed systems is that application-layer threats are not static—they are dynamic—and closed systems are not designed to respond adequately to these types of threats. After their initial configuration, ASIC-based systems cannot be reprogrammed to address new attacks. To combat these new attacks, closed systems include a general-purpose processor that is field programmable; however, the processor lacks the acceleration technology found in ASICs. This results in two major issues for closed systems:

1. Fast and slow inspection tracks: For simple tasks, ASICs enable a fast inspection track that accelerates traffic flow. But for more complex tasks, such as those associated with Web traffic, email, VoIP or accessing sensitive information, traffic is sent to a secondary processor. This slow track is dependent on bus speeds and on processor power. However, the technologies chosen for these components are of secondary concern to the ASICs and often provide underpowered performance against application-layer threats. Because of this, ASIC-based systems are particularly prone to slow performance as new attacks appear, when compared to open systems.
2. Closed systems lose their performance value over time: Since ASICs cannot be programmed in the field to deal with the new threats that appear daily, closed systems start becoming slower the day after the system has been designed. With each additional attack that appears, the closed system will become slower and slower. Over the lifetime of a closed system, it can be expected that the owner will face a choice. Either accept debilitating performance or do not activate any defenses to protect against new attacks.

Security solutions designed with open architectures, based on multi-core high performance processors, deliver the flexibility

and performance needed to protect against existing and potential threats. Clearly not just any CPU can keep up with the most demanding security workloads, but the latest power-optimized Quad-Core Intel® processors are up to the task. By combining the security and platform expertise of Check Point and Intel, institutions can deploy world-class security devices designed to meet next-generation security challenges.

High Security for High-Performance Environments

Security appliances based on Check Point Open Performance Architecture and Quad-Core Intel® Xeon® processors 5400 series are changing the security performance equation. Check Point security software utilizes as many as eight CPU cores, supplied by two quad-core Intel® processors, which provide the performance headroom needed to protect networks into the future. Key performance statistics include:

- 12 gigabits per second (Gbps) firewall inspection delivers data center level speed for the most demanding enterprise environments.
- 5.3 Gbps intrusion prevention inspection with default settings provides a balance between security and performance.
- 1.8 Gbps intrusion prevention inspection with strict protection file offers maximum security without compromising performance.

To reach these speeds, the security appliance employs the Check Point Open Performance Architecture, which consists of three patented technologies:

- CoreXL™ Multi-core Acceleration—distributes security inspection duties throughout all the cores in a multi-core processor-based system, thereby fully utilizing the computing power of the security appliance.
- SecureXL™ Security Acceleration—accelerates security inspection by removing the latency introduced as network traffic passes through a security device.
- ClusterXL™ Smart Load Balancing— provides high availability and load sharing and enables near-linear performance as the cluster size increases. It distributes traffic between clusters of redundant gateways so that the computing capacity of multiple machines may be combined to increase total throughput.

These three technologies work together to fully accelerate security inspection along a unified path that ensures both high performance and high security.

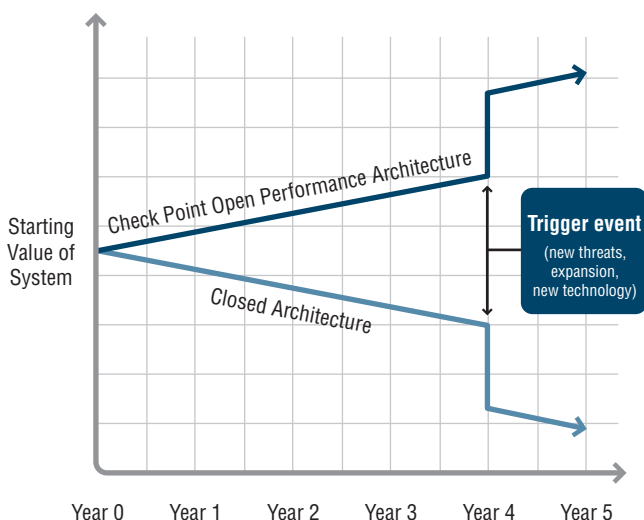


Figure 1. Security Performance Timeline

Increasing Performance Value Over The Solution's Lifetime

The true value of a security system is not whether it can efficiently deal with the threats you know, but rather the threats that have not appeared yet. A key feature of the Check Point Open Performance Architecture is the ability to adapt to new threats while maintaining a predictable level of performance. It does this by maintaining a unified path of inspection based on the use of commonly available hardware technologies. When a new type of attack appears, Check Point security software will not suffer the performance degradation associated with switching from ASIC-based acceleration to general-purpose processors. Rather, the protection will be treated as any existing defense would be processed. This gives customers the knowledge that performance will remain consistent as their security policies adapt.

In comparison, ASIC-based systems start losing value the minute a new attack appears and traffic begins to be transferred to the slower inspection channel. In practice, this results in a steep performance decline when the first deep inspection setting is turned on anywhere from 95 to 99 percent degradation. It is important to understand that this loss of performance value does not start when the system is purchased. It starts when the system is designed. Toward the end of a particular system's life cycle, the value is exceptionally low.

Figure 1 illustrates the concept that two systems—one based on the Check Point Open Performance Architecture and the other a closed architecture—that start with equal performance value will see a gap quickly appear as new attacks emerge. Complicating matters is the fact that every three to four years a new paradigm of security, such as the wide adoption of the Web for business or the emergence of Slammer and Blaster worms, and closed systems

cannot adapt. Because of its open nature, the Check Point solution will be able to provide coverage much faster when these large scale security shifts occur.

Intel Scalable Platforms

Scalability. Power efficiency. Higher performance. Now more than ever, Intel is delivering computing platforms that meet the security challenges of demanding service providers; enterprises; small and medium-sized businesses; and even cost-sensitive consumers, as shown in Figure 2. A cost advantage of using Intel® architecture processors is the same code base can run on a large family of processors: from high-end Quad-Core Intel® Xeon® processors down to value-oriented Intel® Celeron® or Intel® Atom™ processors. This high level of software portability reduces the software development costs for companies, like Check Point, that supply security solutions for a wide range of organizations.

Security systems running on eight CPU cores, supplied by two Quad-Core Intel® Xeon® processors 5400 series, are providing the performance headroom needed to protect networks into the future. Based on the Intel® Core™ microarchitecture, these processors offer breakthrough performance and performance per watt compared to previous-generation single-core processors. This translates into greater performance with fewer cooling challenges and enables security applications to run within a smaller footprint.

The power efficiency of Intel processors can be attributed to advances in processor architecture and chip manufacturing. The latest processors implement new innovative architectural features that save energy, like reducing power to processor circuits that

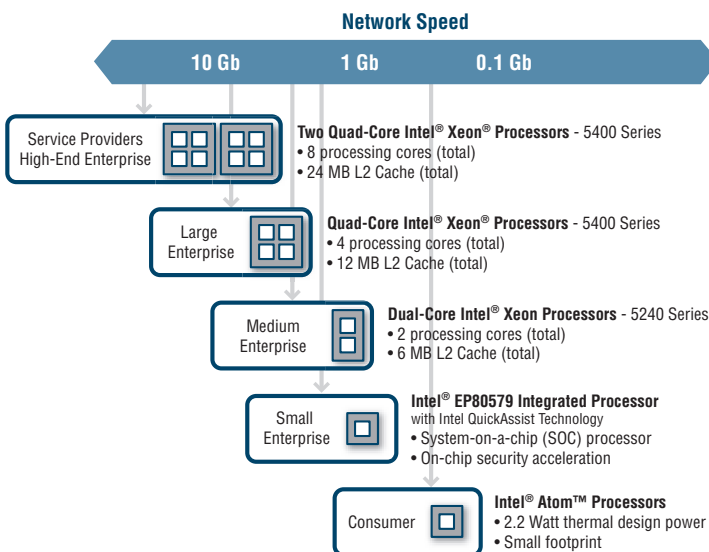


Figure 2. Intel® Processor Scalability Examples

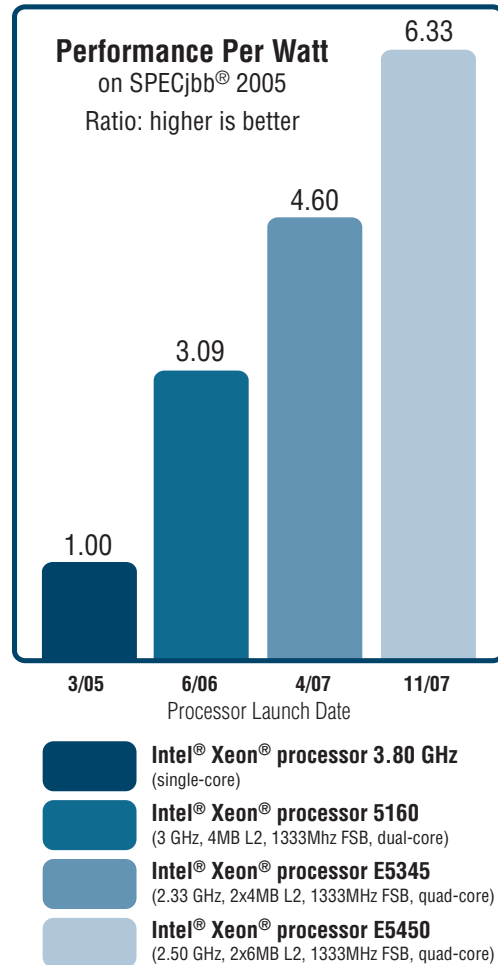


Figure 3. Performance Per Watt Improvements

are idle. Lowering transistor power consumption, the Intel® 45nm technology is expected to fuel the ongoing performance advantages and significantly increase the transistor density over the previous Intel® 65nm technology. The performance per watt advantages are shown in Figure 3, where Intel® quad-core processors have over 600 percent¹ better performance per watt advantages over single core Intel® Xeon processors that were leading-edge just a couple years ago.

Check Point CoreXL: Multi-core Acceleration

Fully utilizing the high compute performance of Intel® multi-core processors, CoreXL introduces advanced load balancing to increase throughput for the deep inspection required to successfully deploy intrusion prevention on the firewall. Providing increased performance for security functions that were previously unaccelerated, CoreXL allows networks to operate at peak performance and security levels.

“Check Point CoreXL running on open multicore architectures provides customers the flexibility to deal with new applications and threats while maintaining a predictably high level of performance. It is the only security solution in the market today that can fully leverage the performance benefits from multicore architectures.”

Dorit Dor

Vice President of Products
Check Point Software Technologies Ltd.

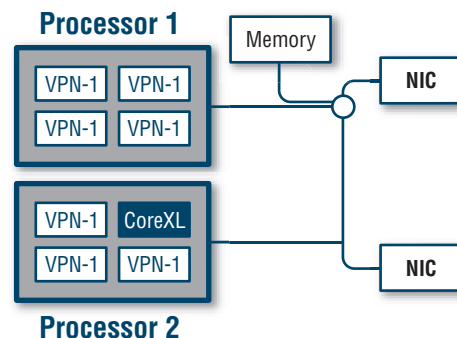


Figure 4. CoreXL manages CPI cores and VPN-1 copies

When the CoreXL technology is activated, it immediately assigns a core to act as a director for traffic. The other cores are designated to run instances of VPN-1. For example, if an appliance contains two quad-core Intel processors, one core will act as the director while the other seven cores run VPN-1, as shown in Figure 4. The core acting as a director has two main functions. First, it decides whether an incoming traffic flow can be accelerated by SecureXL, and then it assigns the flow to an available core, which carries out full security inspection.

CoreXL enables a more efficient distribution of security duties across multiple cores using intelligent load balancing techniques. By balancing the load across multiple cores, Check Point Open Performance Architecture gains a higher level of efficiency than previous multithreaded security applications. These earlier security applications could take advantage of multiple cores by running multiple instances of an application on every core, but could not balance the load equally between them. This created situations where one core would be running at 100 percent utilization while another might be running at 10 percent utilization.

It is the intelligence built into the director core that ensures security processing responsibilities are evenly distributed among all the cores. The result is a 600 percent increase¹ in throughput when CoreXL is activated. The throughput was raised from 300 Mbps to more than 1.8 Gbps. The testing parameters were a strict protection profile with 80 percent of SmartDefense™ settings activated. The network traffic passed through represented a blend of protocols and applications similar to that found on the Internet.**

Securing a Virtual Environment

Today's powerful servers and security appliances often run a mix of applications in a virtualized environment. Virtualization isolates applications, which prevents unintended software interactions that could have dire performance or security consequences. As it's important to protect networks against external threats, security

devices require protection from other applications that may be running on the same platform. This capability is available with Check Point VPN-1 Virtual Edition (VE) running in a VMware® virtual machine, which segregates virtual systems from each other as well as from external threats.

VMware, a leading supplier of virtualization solutions, uses the hardware-assist capability incorporated in Intel processors, called Intel® Virtualization Technology (Intel® VT). Intel VT performs various virtualization tasks in hardware, like memory address translation, which reduce the footprint of the virtualization software and improve the performance of guest operating systems and applications.

Increasing Software Portability

Just as security threats are evolving at a rapid pace, security vendors are quickly adopting the latest technological innovations to keep a step ahead of hackers. Security software and special-purpose security hardware are continuously improving, which is great for network protection, but it creates challenges to software portability. Portability is critical for vendors trying to get solutions to market faster and more reliably.

To facilitate portability, Intel created a software framework that allows vendors to incorporate next-generation security technology,

yet avoid major application software changes. The framework, called Intel® QuickAssist Functional API for Cryptography, enables developers to choose the best platform architecture while keeping the software compatible, irrespective of the technology path. With this portability solution, IT professionals can be ensured of a consistent and effective approach to address security threats across a broad spectrum of products.

Conclusion

The shift from network-layer attacks to dynamically changing application-layer threats has dramatically increased security performance needs. Addressing application-layer threats requires an architecture that can quickly evolve to guarantee performance yet maintain a high level of security. While closed, ASIC-based architectures have not been able to make an efficient shift to protecting against application-layer threats. Check Point Open Performance Architecture and Intel processor and technologies provide the foundation needed by telecom operators, large campuses and data centers to gain high performance while maintaining a high level of security.

With this open architecture, organizations can deploy security that delivers on the promise of integrated intrusion prevention without fearing the loss of network performance.



www.intel.com



Check Point
SOFTWARE TECHNOLOGIES LTD.

www.checkpoint.com

¹ Performance tests and ratings are measured using specific computer systems and/or components and reflect approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, visit http://www.intel.com/performance/resources/benchmark_limitations.htm

** Results based on Check Point internal testing, April 2008

Copyright © 2008, Intel Corporation. All rights reserved. Intel, the Intel logo, Atom, Celeron and Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States or other countries.

* Other names and brands may be claimed as the property of others. Printed in USA MS/SD/1108 Please Recycle Order No. 320923-001US