

White Paper

Designing Firewall/VPN with the PowerQUICC™ III MPC8572E

Overview

This white paper is the first of a series of white papers on using the MPC8572E to design different types of security equipment including firewall/Virtual Private Network (VPN), Intrusion Detection Systems (IDS)/Intrusion Prevention Systems (IPS), anti-virus/anti-spam/content filtering and Unified Threat Management (UTM)/Integrated Services Router (ISR).

Contents

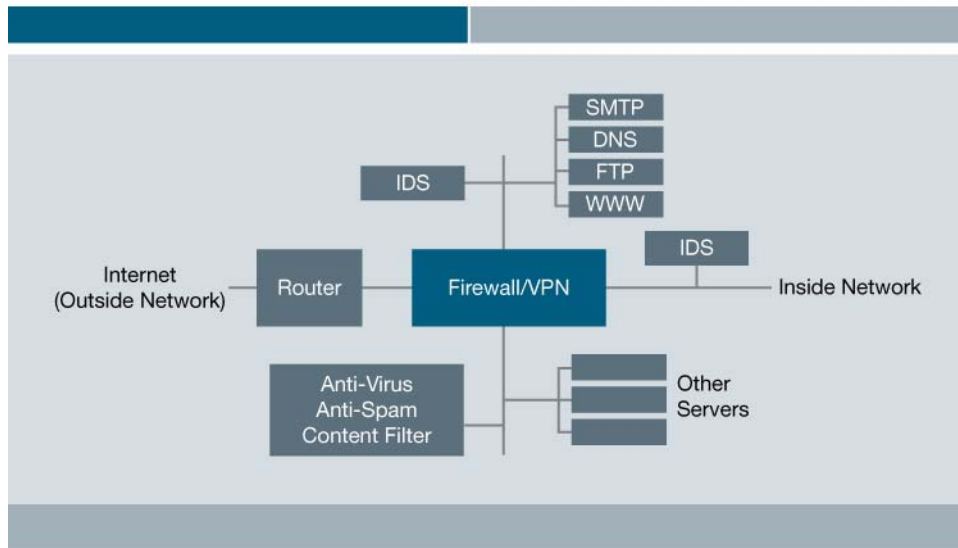
1	Network Security Equipment Overview and Trend	1
2	Vendor Pain Points and Design Challenges	3
3	MPC8572E PowerQUICC™ III Processor Overview	4
4	Designing Firewall/VPN with the MPC8572E.....	5
4.1	System Design with the MPC8572E.....	5
4.2	Firewall.....	5
4.2.1	Information Flow	5
4.2.2	Flow-Based Firewall Operations.....	6
4.2.3	Firewall Operations on the MPC8572E.....	7
4.2.4	Performance Advantages of the MPC8572E in Firewall Data Path.....	8
4.3	IPSec VPN.....	9
4.3.1	IPSec VPN Operations.....	9
4.3.2	IPSec VPN Operations on the MPC8572E.....	10
4.3.3	Performance Advantages of the MPC8572E in IPSec VPN	11
4.4	Anticipated Features	11
5	Summary	12



1 Network Security Equipment Overview and Trend

The stateful inspection firewall/IPSec VPN security gateway has been—and still is—the most critical piece of network security equipment for most enterprises. Firewall/VPN is a perimeter-defense device, typically deployed where the enterprise's internal network meets the open Internet. The main purpose of the firewall is to stop unwanted traffic from entering or leaving the internal enterprise network. The purpose of the IPSec VPN is to provide secure communication between two sites through the open Internet.

Firewall/VPN Deployment



While the firewall/VPN security gateway is extremely important, some 90 percent of attacks in recent years have exploited application vulnerabilities. The traditional stateful inspection firewall, based largely on matching packet header information against Access Control Lists (ACLs), is ineffective to fend off such attacks.


An Intrusion Detection System (IDS) will expose application layer attacks. Some enterprises deploy an IDS to monitor traffic in vital network segments. But detection alone is insufficient—it is also important to terminate the attack upon detection. The trend is to evolve the IDS into an Intrusion Prevention System (IPS), which takes detection to the next level and stops the detected attacks, including application attacks.

In some instances, application-specific anti-virus, anti-spam and content filtering devices are deployed in enterprise networks to complement the firewall/VPN.

One undesirable outcome of this trend has been the proliferation of network security devices, which increases cost and management complexities. There are too many pieces of equipment to buy and operate. Some IT managers are seeking simpler solutions with Universal Threat Management (UTM) systems or Integrated Services Routers (ISR) that incorporate multiple networking and security functions, such as routing, firewall, IPSec, IDS/IPS, anti-virus, anti-spam and content filtering into a single device.

¹Symantec Internet Security Threat Report, Trends for July 05–December 05, Volume IX, March 2006

Trend Towards UTM/ISR



Networks—both Wide Area and Local Area—are speeding up. Network users expect faster and faster response time. IT managers do not want networking devices to bottleneck traffic and realize that additional security measures can do just that. To be acceptable solutions, the speed of perimeter network security devices must ramp at the same speed as the Wide Area Network.

Network designers and managers are increasingly taking the view that perimeter defense alone is inadequate. The use of Portable PCs and wireless access increase the probability of “attacks from within.” To prevent the spread of worms or viruses from both outside and inside, network security devices should be deployed within the enterprise network, or even incorporated into the LAN infrastructure. LAN speed is typically higher than WAN speed, so network security devices deployed in the LAN network demand even higher performance.

In summary, there is a definite trend in network security towards:

- Application content security
- Higher integration
- Higher speed

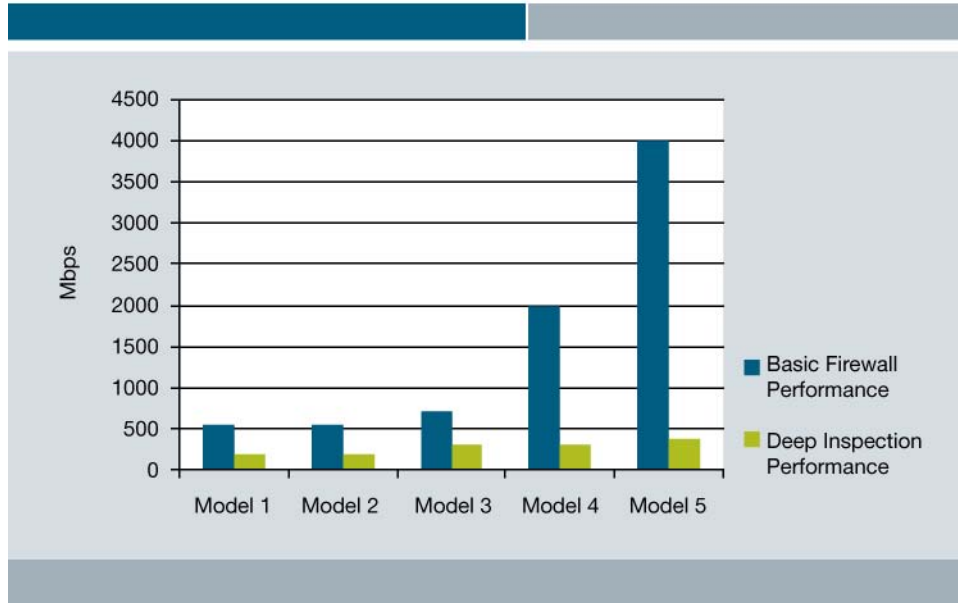
Despite the trend towards application content security and higher integration, firewall/VPN remains the most important piece of network security equipment in most enterprises. For enterprises deploying UTM or ISR instead, firewall/VPN capabilities are the foundation of these devices, upon which additional application content security capabilities can be added.

2 Vendor Pain Points and Design Challenges

The traditional firewall/VPN has been around for quite a few years. Interestingly, the driving force behind the development of next-generation firewall/VPN solutions for most vendors is not just continuous improvement in cost and performance—it is the perceived limitations in current generation firewall/VPN solutions to protect against application-layer attacks. Even if the new development was triggered by reasons other than security, R&D investment in a new firewall/VPN without taking the content security trend into account is myopic.

Processing application content is very CPU-intensive. The graph below shows that using software to perform “deep packet inspection” does not lead to the desired line-rate speed.

Firewall Performance: Basic vs. “Deep Inspection”



Given the continually changing environment and the focus on content security, higher integration and higher speed, some vendors are considering designing a new, single, hardware platform with high packet and content processing capabilities to maximize R&D effectiveness and minimize time to market. Depending on the software loaded, this platform can perform as a firewall/VPN, IDS/IPS, anti-virus/anti-spam/content filter or ISR/UTM. And for all the different network security products, the performance must remain high and the cost must be maintained.

To achieve this goal,

- The processor(s) powering the platform has to have the right capability and performance at the right price point.
- The system design has to be simple.
- The software architecture for various products using the same platform has to be consistent.
- There needs to be a significant re-use of firewall/VPN software.

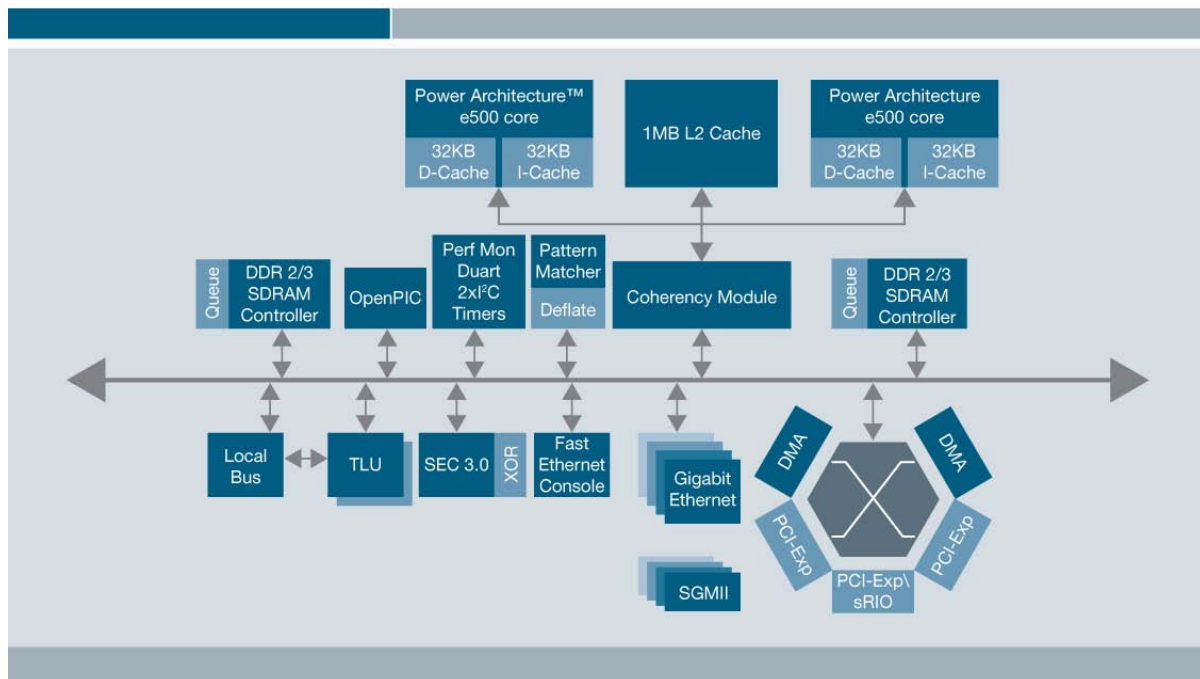
3 MPC8572E PowerQUICC™ III Processor Overview

Freescale's next-generation PowerQUICC™ III integrated communications processors are designed to provide solutions for symmetrical and asymmetrical multi-core systems. Based on the scalable e500 system-on-chip (SoC) platform built on Power Architecture™ technology, they deliver dual-core gigahertz-plus communications processing performance with advanced content processing and security features.

The MPC8572E family of processors is designed to offer clock speeds from 1.2 GHz to 1.5 GHz, combining two powerful processor cores, enhanced peripherals and high-speed interconnect technology to balance processor performance with I/O system throughput. These processors also contain an application acceleration block that integrates four powerful engines: a table look-up unit (TLU) that offloads complex table searches and header inspections; a pattern-matching engine to handle regular expression matching; a deflate engine to manage file decompression; and a security engine that accelerates crypto operations in IPsec and SSL/TLS for virtual private networks.

Based on Freescale's 90 nm silicon-on-insulator (SOI) copper interconnect process technology, the MPC8572E is designed to deliver higher performance with lower power dissipation. The MPC8572E processors provide a significant performance increase and represent the next step in continuous innovation from the popular PowerQUICC family. With unmatched integration, the MPC8572E platform builds on the embedded core performance of Power Architecture technology and adds new features to enhance traffic management and security acceleration. Support for high-speed interfaces on the MPC8572E enables scalable connectivity to network processors and/or ASICs in the data plane while the MPC8572E platform handles complex, computationally demanding control plane processing tasks. These processors also include a next-generation double data rate (DDR2/DDR3) memory controller, enhanced Gigabit Ethernet support, double precision floating point and an integrated security engine that features updated Advanced Encryption Standard (AES) functionality.

MPC8572E



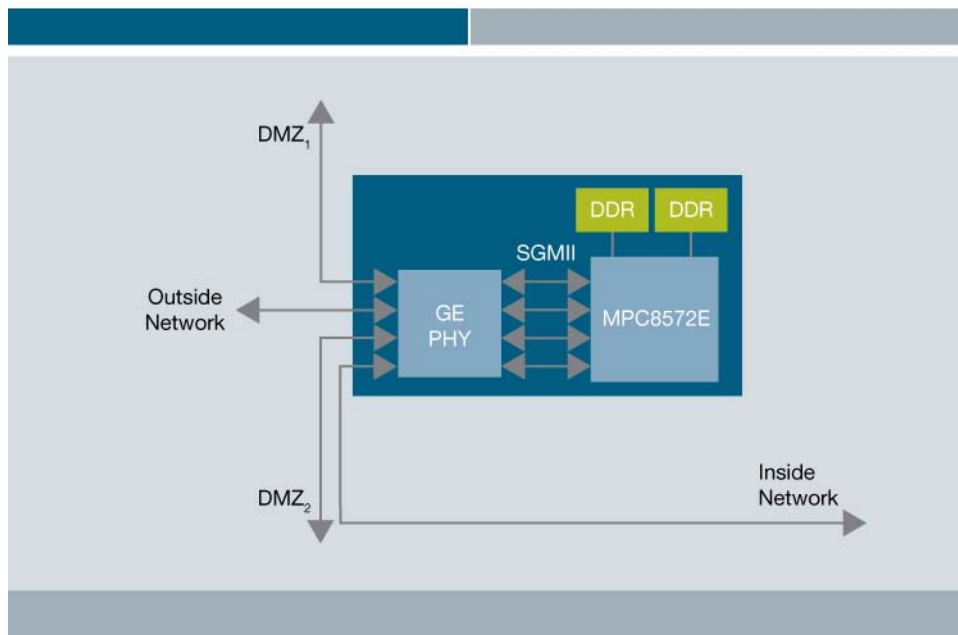
4 Designing Firewall/VPN with the MPC8572E

4.1 System Design with the MPC8572E

The simplified block diagram below shows the essence of a 4-port security appliance, illustrating how easy a system design using the MPC8572E can be. The very simple system design—a direct result of the exceptional integration in the MPC8572E—enables significantly lower system cost and shorter time to market.

Subsequent sections will describe in more detail how all major processing and I/O elements are integrated directly into the MPC8572E, with a highly optimized internal interconnect architecture to ensure high bandwidth, low latency and efficient pipeline operation. As a result, the MPC8572E enables high-performance as well as cost-effective system design.

4-port Security Appliance



4.2 Firewall

4.2.1 Information Flow

To understand firewall operation, let's examine the information flows of two very well-known service transactions:

- World Wide Web (HTTP)
- File Transfer (FTP)

The flow of packets from a client to a server in World Wide Web service using the HTTP protocol is characterized by the selector values in the following 5-tuple:

<protocol, SA, DA, SP, DP>, where

- Protocol = TCP
- SA = address of the client
- DA = (typically published) address of the server
- SP = port number of the client
- DP = 80, IANA-assigned port number for HTTP protocol

In the other direction of the flow, the selector values between the source and destination are swapped.

FTP file transfer consists of a connection for control and a connection for data transfer. The client-to-server control flow is characterized by the selector values in the following 5-tuple:

<protocol, SA, DA, SP, DP>, where

- Protocol = TCP
- SA = address of the client
- DA = (typically published) address of the server
- SP = address of the client
- DP = 21, IANA-assigned port number for FTP control

The data connection is a child flow of the control connection, which is the parent flow. The client uses the PORT or the PASV commands to specify the address and port number to be used to transfer the data requested. The FTP server-to-client data flow is characterized by the selector values in the following 5-tuple:

<protocol, SA, DA, SP, DP> where

- Protocol = TCP
- SA = (typically published) address of the server
- DA = dynamic address specified in the PORT or PASV command
- SP = typically dynamic port number assigned by the server (or the default value of 20, IANA-assigned port number for FTP data)
- DP = dynamic port number specified in the PORT or PASV command

The previous examples illustrate that information services are flow-based. The flows are characterized by their 5-tuple selectors: <protocol, SA, DA, SP, DP>. The selector values do not change during the lifetime of a flow.

The address of the server is typically published. The port number indicating the service is IANA-defined. A single flow characterized by well-known server address and port number is applicable to many, but not all, services.

There are multimedia services that use more than one child flow. The parent (typically control) flow negotiates the selector values in the child (typically data or media) flows.

4.2.2 Flow-Based Firewall Operations

Well-known port numbers enable relatively simple policy enforcement operation based on an Access Control List (ACL). For example, for the case of Web traffic, the applicable ACL can be defined as follows:

“Allow traffic from the Outside Network to DMZ1 if the protocol is TCP and the destination port number is 80.”

The policy enforcement stage involves matching the selector value of a received packet against ACLs configured. Once there is a match, all subsequent packets belonging to the flow is “allowed”, or simply forwarded in this case.

For the FTP case, policy enforcement is more involved. In addition to matching the applicable ACL, which will look similar to:

“Further process traffic from the Outside Network to DMZ1 if the protocol is TCP and the port number is 20.”

The control flow packets are “further processed” in that the FTP protocol exchange is followed to determine the selector values of the child data flow. A pinhole corresponding to the child flow’s negotiated selector values is opened up. Subsequent packets matching the pinhole selector values are allowed to go through.

Note that in general:

1. Policy enforcement is more complex and time-consuming than data transfer.
2. Once policy enforcement is performed at the beginning of a flow, it does not need to be repeated for subsequent packets in the same flow or its child flows.

Firewall designers achieve high performance by exploiting these characteristics of information flows. The typical algorithm is:

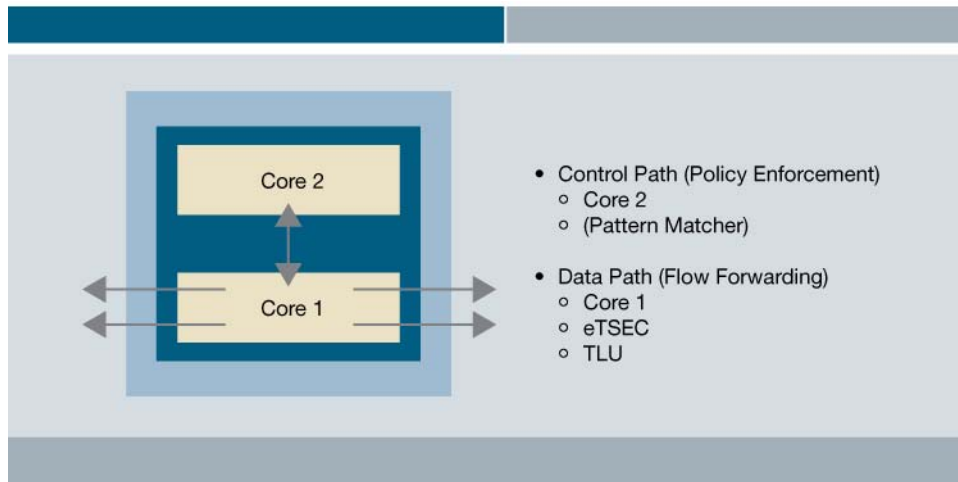
When a packet is received, its selector values are matched against a flow table of established flows.

- No match – execute control path (policy enforcement). Match relevant received packet header fields against ACL configured and scrutinize application protocol and content if appropriate. Based on the policy, the methods of processing the packet's own flow and its child flows, if any, are determined and entered into the flow table. The received packet is processed accordingly.
- Match – execute data path. Process received packet, e.g., drop, forward, NAT, tunnel ... according to the recipe recorded in the flow entry.

The control path determines the performance of the first or first few packets and the data path determines the performance of the rest of the packets in the flow. Since most information flows involve more than a few packets, the data path largely determines the throughput performance of the firewall. Firewall designers therefore highly optimize the data path in order to achieve high system performance.

4.2.3 Firewall Operations on the MPC8572E

Dual-Core Usage Model for Firewall Operations



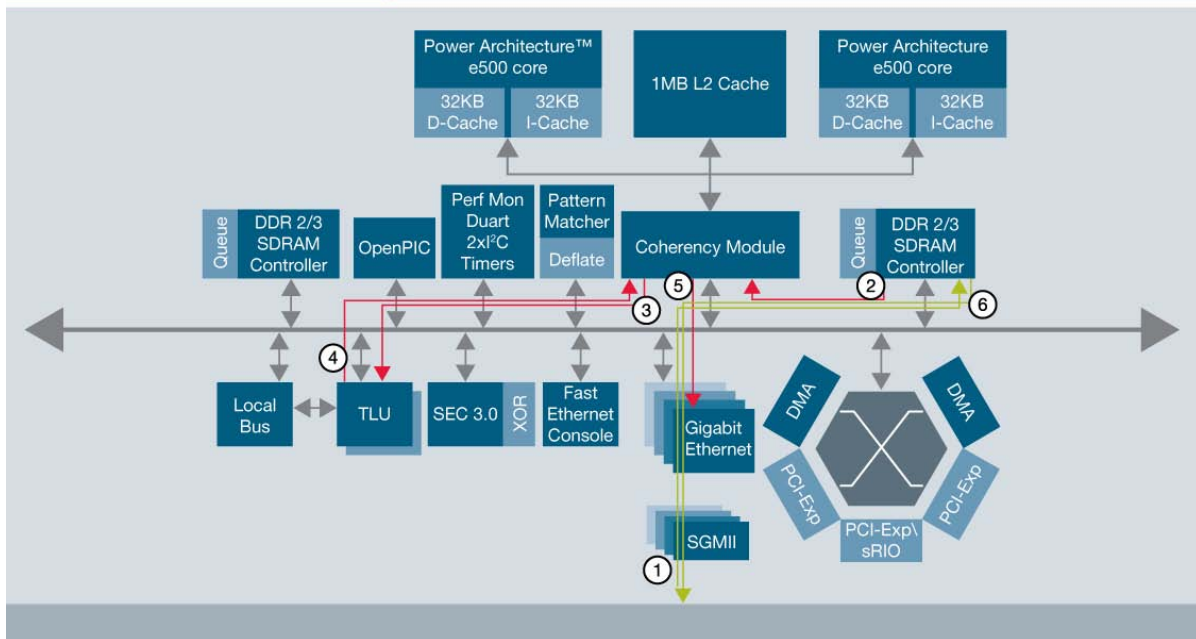
Firewall operations are well-suited to the Asymmetric Multi-Processing (AMP) usage model on the dual-core MPC8572E:

- Control Path (Policy Enforcement) runs on dedicated CPU Core2, possibly with the help of the Pattern Matcher for parsing application protocol and content.
- Data Path runs on CPU Core1 with the TLU used to accelerate flow table lookup and eTSEC, network I/O.

The firewall data path is shown in the diagram on the following page:

1. eTSEC puts received packet into memory and interrupts Core1
2. Core1 extracts 5-tuple selector from the packet header
3. Core1 writes selector to the TLU to lookup flow table
4. Core1 reads back (+ve) results of the lookup and retrieves additional flow entry data from memory (if required), increments statistics, adjusts header and formats the packet for transmission
5. Core1 instructs appropriate eTSEC to transmit packet
6. eTSEC transmits packet

Firewall Data Path on MPC8572E



4.2.4 Performance Advantages of the MPC8572E in Firewall Data Path

4.2.4.1 Packet I/O

The firewall is first and foremost a networking device. As such, it should be able to receive and transmit packets at a high rate as a prerequisite. The overhead in servicing a high rate of transmit and receive interrupts is high and can significantly slow down the performance of the firewall. The integrated eTSEC (Enhanced Triple Speed Ethernet Controller) is able to coalesce interrupts, thereby reducing the interrupt servicing overhead and improve performance.

In a firewall where the I/O rate is high and computation logic relatively low, it is often memory access speed and not the availability of CPU cycles that determine system performance. In the MPC8572E, the eTSEC stashes received packet headers in L2 cache while writing to memory. As a result, the e500 CPU core accesses data with reduced latency. In fact, the transmit and the receive buffer descriptors can be locked in the L2 cache for fast access by the eTSEC and the e500 core.

4.2.4.2 Flow Table Lookup

As described earlier in Section 4.2.2, searching for an existing entry in a potentially very large flow table is performed every time a packet is received. The MPC8572E's built-in TLU provides hardware acceleration to this operation.

4.2.4.3 Data Path Processing

In order to provide high performance, the firewall data path is highly optimized to minimize the number of instructions. The eTSEC helps in offloading IP and TCP checksum calculations.

4.3 IPSec VPN

4.3.1 IPSec VPN Operations

IPSec VPN operations can be separated into two distinct parts:

- Internet Key Exchange (IKE)
- IPSec data path

The purpose of the first part is to set up Security Association (SA) for IPSec data path. Public key techniques or pre-shared key is used to mutually authenticate the communicating parties. Diffie-Hellman key exchange is used to set up a shared session secret from which cryptographic keys for IPSec are derived.

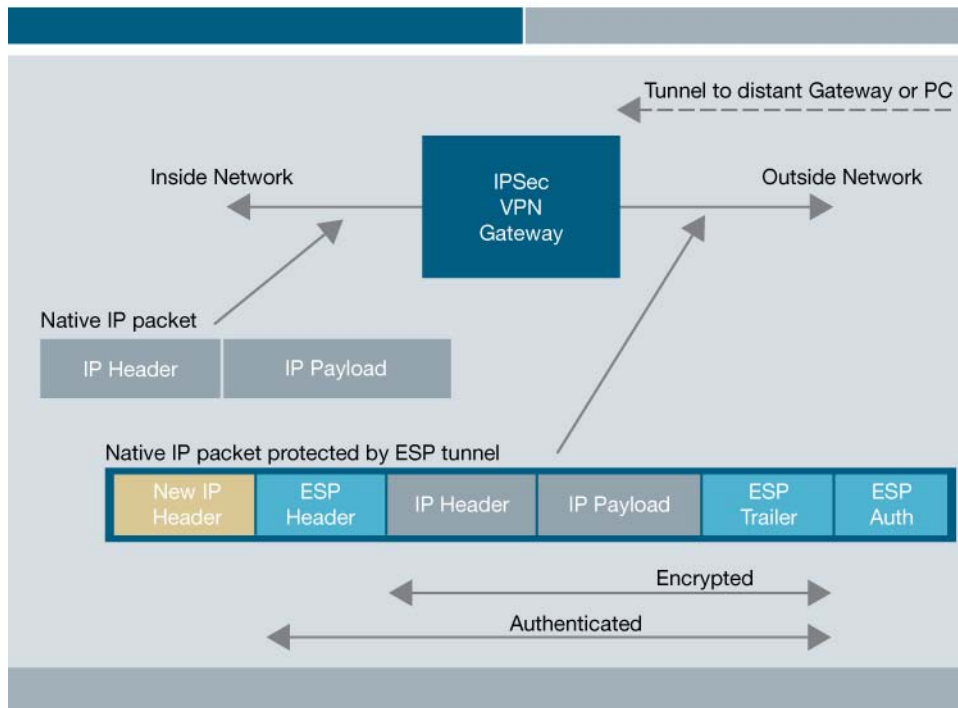
IKE operation is highly complex in terms of cryptographic operation and protocol. This part of the VPN operations determines VPN tunnel setup time.

In the IPSec data path, the two parties communicate using the negotiated Security Association, e.g.:

- IPSec protocol = ESP
- Encrypt Algorithm = 3DES
- Hash Algorithm = MD5
- Encapsulation = tunnel
- ...

The following diagram illustrates the end result of typical IPSec data path operation—entering an encrypted tunnel.

Example IPSec Operation: Packet Transformation



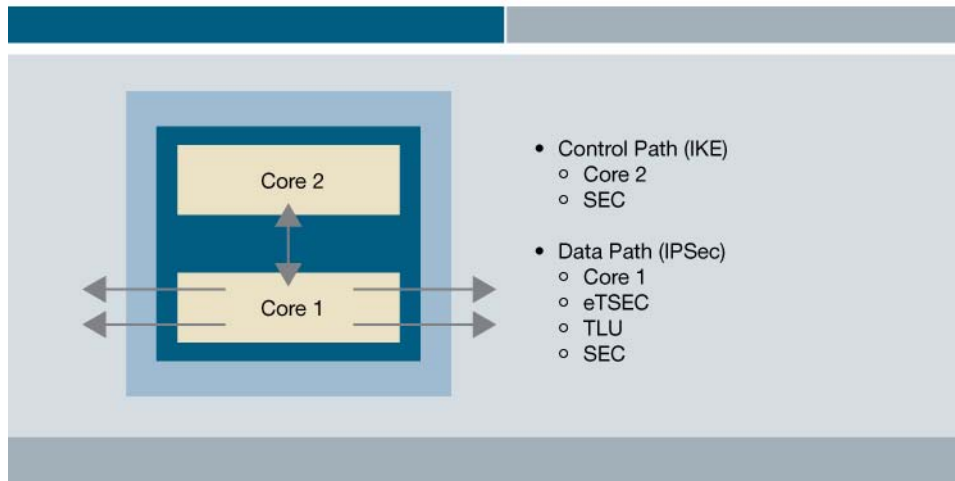
The IPSec data path protocol is relatively simple, but the cryptographic operation is complex. The data path part largely determines the throughput performance of the system and therefore must be highly optimized.

4.3.2 IPsec VPN Operations on the MPC8572E

IPsec VPN Operation is well-supported in the MPC8572E using the Asymmetric Multi-Processing (AMP) usage model:

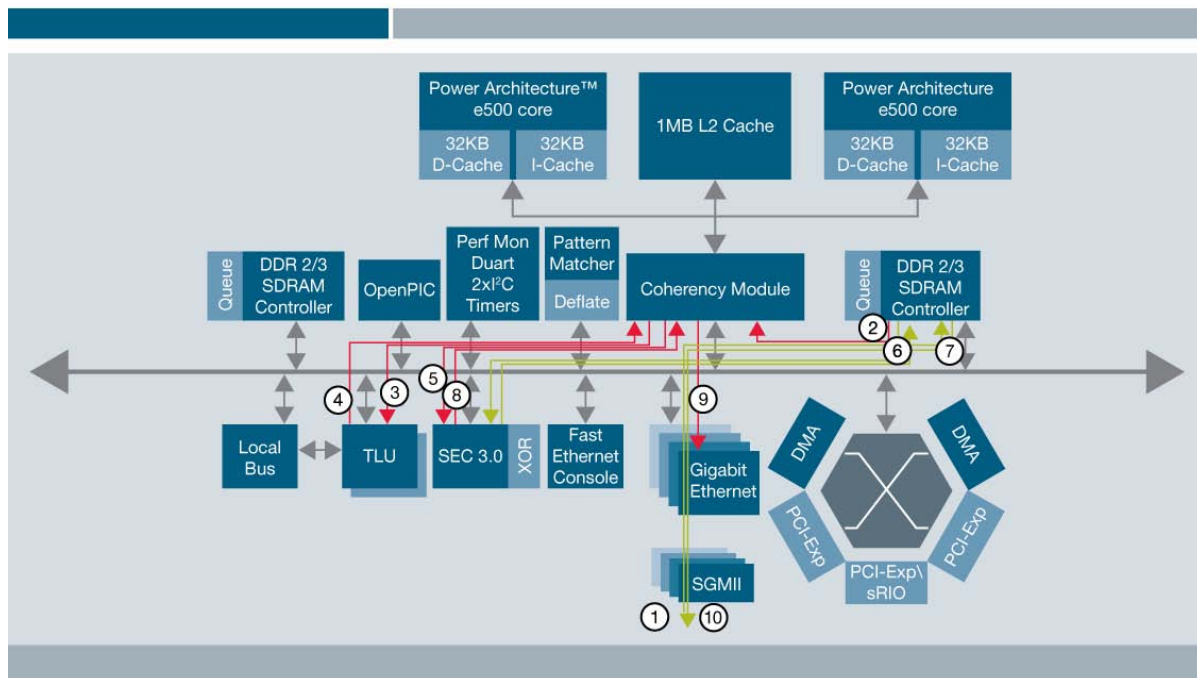
- Control Path (IKE) runs on the e500 Core2, with cryptographic operations offloaded to the SEC block.
- Data Path (IPsec) is orchestrated by the e500 Core1, with eTSEC(s) enabling and accelerating packet I/O, TLU speeding up Security Policy Database and Security Association Database lookup, and the SEC block accelerating the cryptographic operations.

Dual-Core Usage Model for IPsec Operations



Operation of IPsec data path entering the encrypted tunnel is described in more detail in the following diagram:

IPSEC Data Path on MPC8572E



1. eTSEC puts received clear-text packet from “inside network” into memory and interrupts Core1
2. Core1 extracts 5-tuple selector from the packet header
3. Core1 writes selector to the TLU to lookup Security Policy Database
4. Core1 reads back lookup results and follows pointer to Security Association Database entry to retrieve key material
5. Core1 informs SEC of location of data and key material
6. SEC reads clear data from memory
7. SEC writes encrypted data to memory
8. Core1 reads results from SEC
9. Core1 formats encrypted data and informs “outside network” eTSEC
10. Encrypted Packet is transmitted to “outside network”

4.3.3 Performance Advantages of the MPC8572E in IPSec VPN

The performance advantages of the MPC8572E in firewall data path described in Section 4.2.4 also apply to the IPSec VPN. This includes packet I/O, table lookup and minimizing data path processing by the eTSEC.

Specific to IPSec VPN, the IKE stage requires highly complex public key computation. The integrated SEC block accelerates this operation, thereby improving tunnel setup time.

The IPSec data path also requires complex cryptographic operations, which are offloaded and accelerated by the integrated SEC block. The interaction between the e500 core and the SEC block is very efficient—the two parties communicate efficiently via crypto descriptors in L2 cache and the IPSec transformation is performed in a single pass.

4.4 Anticipated Features

This white paper focuses on the design of traditional firewall/VPN. It was discussed earlier in Section 2 that any new network security equipment design should take the anticipated requirement of content security into account, because 90 percent of successful attacks are caused by application-level vulnerabilities.

In the context of the HTTP and FTP examples in Section 4.2.1, the HTTP or FTP requests may be a buffer overflow, directory traversal or other kind of application attack, even though the ACL says that the request is allowed. Also, the objects being retrieved may contain virus, confidential material or other undesirable content that should be stopped.

While these content security features may not be needed on day one, it is a competitive advantage to be able to support these features and at high speed when they are required on the same platform, provided the cost associated with this latent capability is low.

The MPC8572E contains hardware acceleration blocks for pattern matching and decompression, which are very useful in speeding up content security operations. And, since these blocks are integrated in the SoC instead of separate devices, the system cost can be kept low.

5 Summary

The MPC8572E is a PowerQUICC III processor optimized for network security processing at Gbps speed. Its dual e500 cores provide CPU cycles and flexibility for performing various security operations. The MPC8572E's integrated hardware blocks (Pattern Matcher, Deflate, Table Lookup Unit and Security Engine) further accelerate packet header and application content processing with low power dissipation, while its highly optimized internal interconnect architecture ensures high bandwidth, low latency and efficient pipeline operation.

The MPC8572E processor—with all major processing and I/O elements included—enables very simple, elegant platform designs, with low system cost and a short design cycle. Because it evolves from the proven PowerQUICC III SoCs implemented with the same e500 core, eTSEC and SEC blocks, it makes significant reuse of existing firewall/VPN possible.

While this white paper focuses on describing how the MPC8572E can be used for firewall/VPN design, complementary white papers will show that the SoC can also be used for other security devices, including IDS/IPS, anti-virus/anti-spam and ISR/UTM, all with a consistent architecture. In other words, a vendor can design a single hardware platform with the MPC8572E and turn it into a specific network security device with the appropriate software load at the appropriate time.





How to Reach Us:

Home Page:

www.freescale.com

e-Mail:

support@freescale.com

USA/Europe or Locations Not Listed:

Freescale Semiconductor
Technical Information Center, CH370
1300 N. Alma School Road
Chandler, Arizona 85224
1-800-521-6274
480-768-2130
support@freescale.com

Europe, Middle East and Africa:

Freescale Halbleiter Deutschland GmbH
Technical Information Center
Schatzbogen 7
81829 Muenchen, Germany
+44 1296 380 456 (English)
+46 8 52200080 (English)
+49 89 92103 559 (German)
+33 1 69 35 48 48 (French)
support@freescale.com

Japan:

Freescale Semiconductor Japan Ltd.
Headquarters
ARCO Tower 15F
1-8-1, Shimo-Meguro, Meguro-ku,
Tokyo 153-0064, Japan
0120 191014
+81 3 5437 9125
support.japan@freescale.com

Asia/Pacific:

Freescale Semiconductor Hong Kong Ltd
Technical Information Center
2 Dai King Street
Tai Po Industrial Estate,
Tai Po, N.T., Hong Kong
+800 2666 8080
support.asia@freescale.com

For Literature Requests Only:

Freescale Semiconductor
Literature Distribution Center
P.O. Box 5405
Denver, Colorado 80217
1-800-441-2447
303-675-2140
Fax: 303-675-2150
LDCForFreescaleSemiconductor@hibbertgroup.com

Information in this document is provided solely to enable system and software implementers to use Freescale Semiconductor products. There are no express or implied copyright license granted hereunder to design or fabricate any integrated circuits or integrated circuits based on the information in this document.

Freescale Semiconductor reserves the right to make changes without further notice to any products herein. Freescale Semiconductor makes no warranty, representation or guarantee regarding the suitability of its products for any particular purpose, nor does Freescale Semiconductor assume any liability arising out of the application or use of any product or circuit, and specifically disclaims any and all liability, including without limitation consequential or incidental damages. "Typical" parameters which may be provided in Freescale Semiconductor data sheets and/or specifications can and do vary in different applications and actual performance may vary over time. All operating parameters, including "Typicals" must be validated for each customer application by customer's technical experts. Freescale Semiconductor does not convey any license under its patent rights nor the rights of others. Freescale Semiconductor products are not designed, intended, or authorized for use as components in systems intended for surgical implant into the body, or other applications intended to support or sustain life, or for any other application in which the failure of the Freescale Semiconductor product could create a situation where personal injury or death may occur. Should Buyer purchase or use Freescale Semiconductor products for any such unintended or unauthorized application, Buyer shall indemnify and hold Freescale Semiconductor and its officers, employees, subsidiaries, affiliates, and distributors harmless against all claims, costs, damages, and expenses, and reasonable attorney fees arising out of, directly or indirectly, any claim of personal injury or death associated with such unintended or unauthorized use, even if such claim alleges that Freescale Semiconductor was negligent regarding the design or manufacture of the part.

Freescale™ and the Freescale logo are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. ARM is the registered trademark of ARM Limited. ARM9, ARM11, and ARML210™ are the trademarks of ARM Limited. Java and all other Java-based marks are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries. The Power Architecture and Power.org word marks and the Power and Power.org logos and related marks and service marks licensed by Power.org.
© Freescale Semiconductor, Inc. 2007.

Document Number: FIREWALLVPNWP
REV 0

