

White Paper

Designing IDS/IPS with the PowerQUICC[®] III MPC8572E

Overview

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

Contents

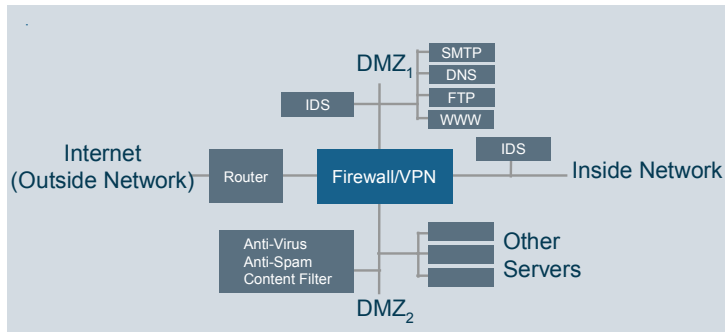
1	IDS/IPS and network security equipment overview and trend.....	3
2	Vendor pain points and design challenges	4
3	MPC8572E PowerQUICC® III processor overview.....	6
3.1	Pattern Matcher	7
4	Designing IDS/IPS with the MPC8572E.....	7
4.1	IDS/IPS processing.....	7
4.2	IDS/IPS operations on the MPC8572E.....	8
4.3	Performance advantages.....	9
4.3.1	Packet I/O	9
4.3.2	Packet processing	9
4.3.3	Application content processing	9
4.3.4	Stateful pattern matching	10
4.4	Accuracy (with performance) advantages	11
4.4.1	Regex	11
4.4.2	Stateful Rule.....	11
4.4.3	Set and subset.....	11
4.4.4	Matching across packet boundaries	11
4.4.5	Performance minimally dependent on number of signatures.....	11
4.5	Hardware platform design with the MPC8572E	12
4.5.1	Cost advantages.....	12
5	Summary	12

1 IDS/IPS and Network Security Equipment Overview and Trend

Next to the Firewall/VPN, the network Intrusion Detection System (IDS) has become an important part of the enterprise's Internet security defense arsenal.

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. The IDS is traditionally deployed to monitor traffic in vital segments in the network, generating alerts when an intrusion is detected.

Figure 1: IDS/IPS and Network Security Equipment



The importance of the IDS has grown significantly as the industry recognizes that 90 percent of attacks in recent years have exploited application vulnerabilities.¹ The traditional stateful inspection firewall, based largely on matching packet header information against Access Control Lists (ACLs), is ineffective to fend off such attacks. A good IDS, on the other hand, can expose these application layer attacks.

But detection alone is insufficient—it is also important to terminate the attack upon detection. Hence, 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 addition to the IDS/IPS, application content security arsenal in an enterprise may also include anti-virus, anti-spam and content filtering devices.

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. Enforcing consistent security policy across multiple systems is inherently problematic.

Some IT managers are seeking simpler solutions with Universal Threat Management (UTM) systems or Integrated Services Routers (ISRs) 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 (Figure 2).

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

Figure 2: Trend Towards UTM



The performance requirement imposed on network security devices—the IDS/IPS in particular—is getting more stringent. Networks—both Wide Area and Local Area—are speeding up. To be acceptable solutions, the speed of perimeter network security devices must ramp at the same speeds as the Wide Area Network. IDS/IPS is not only deployed at the perimeter, it is often deployed to monitor/control traffic between LAN segments, typically with speeds higher than that of WAN. Furthermore, to save cost, IT managers want to use a single IDS/IPS to monitor/control traffic on multiple LAN segments.

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

- Application content security
- Higher integration
- Higher speed

The trend towards application content security leads to the increased importance of the IDS, which in turn is metamorphosing to the IPS. For enterprises deploying UTM or ISR instead of individual network security devices, the IDS/IPS function is an integral and vital component.

2 Vendor Pain Points and Design Challenges

The biggest challenges to the IDS/IPS designer are most likely caused by the migration from detection to prevention. With the trend towards IPS, any new network intrusion system being developed will most certainly be designed to operate in both the detection (monitor traffic only) and prevention (monitor and control traffic) modes.

Superficially, developing an IPS from an established IDS base looks almost trivial in concept—just add a few lines of code to let packets that have been received and checked out to be non-malicious pass through. The reality is that adding the forwarding function is the easy part. The major challenge is to meet the significantly more stringent requirement associated with a network security device operating in-line in terms of:

- System performance
- Detection accuracy

As an IDS—monitor only—missing a packet occasionally is certainly not good, but arguably not quite the end of the world either. On the other hand, as an integral element of the network operating in-line, the IPS is expected to have line-rate throughput, and not adding any appreciable delay. Very few end-users and IT managers are willing to sacrifice network performance for security. An IPS that slows down normal traffic is simply not acceptable.

Compounding this stringent performance requirement are three different factors:


1. the trend towards higher WAN and LAN speed
2. the deployment of IPS between internal LAN segments in addition to being a perimeter defense device between the internal LAN and external open Internet
3. the intention of IT managers to monitor/protect multiple segments with one system to save cost

In addition to high performance, accuracy is another important – perhaps even more important—factor. “False negative” (unable to detect a real attack) and “false positive” (mistakably reporting an attack when there isn’t one) are certainly not good in an IDS. The direct impact is on the analyst in the IT department whose job is to analyze the alerts coming from the IDS. For an IPS, the impact of false positive is much bigger – it means good, normal user traffic is stopped – resulting in irate users, loss revenue and bad company image. Certainly, this behavior of dropping good traffic will not be tolerated. In fact, the wide-spread deployment of the IPS depends on its accuracy.

Unfortunately, IT managers are still put into an awkward position of choosing between performance and accuracy. A later chapter will describe that most IDS/IPS depend on matching of network data against attack signatures. For most systems, the performance decreases with the number of signatures configured. In order to catch more attacks, more signatures have to be configured. In order to achieve high performance, the number of configured signatures has to be reduced.

Today, most Intrusion Detection/Prevention Systems are made simply by loading an off-the-shelf server with appropriate software. But the demand for Gbps speed—and high accuracy—has pushed the pure software approach over the edge. Processing application content to detect application layer attacks as required in the IDS/IPS is very CPU-intensive. Figure 3 shows that using software to perform “deep packet inspection” results in a significant drop in speed.

Figure 3: Firewall Performance: Basic vs. “Deep Inspection”



Given the requirement for Gbps, in-line, application content processing required in high-end IPS, to compete effectively, the IDS/IPS vendor needs a platform that can provide:

- high performance beyond that of a standard CPU – a platform that is optimized for the packet and application content processing operations in the IDS/IPS
- high accuracy without draining the CPU of its precious processing cycles
- competitive product and development cost

Furthermore, particularly for ISR/UTM vendors, the platform must enable an architecture in which various functions including routing, firewall/VPN, IDS/IPS, anti-virus/anti-spam/content filter can operate synergistically.

² Ptacek and Newsham, “Insertion, Evasion, and Denial of Service: Eluding Network Intrusion Detection”, 1998.

3 MPC8572E PowerQUICC® III Processor Overview

The MPC8572E is a new PowerQUICC® III processor purposely built to meet the requirements of high-performance application-aware networking and content security. It is based on the highly successful PowerQUICC system-on-chip (SoC) platform, well-proven in traditional networking, and enhanced with further integration of new hardware, optimized to process application content at high speeds.

The MPC8572E consists of dual e500 cores built on Power Architecture™ technology, achieving clock speeds from 1.2 GHz to 1.5 GHz. The CPU cores, each with 32 KB I-Cache and 32 KB D-Cache, share 1024 KB of integrated L2 cache. For memory, the MPC8572E includes two integrated 64-bit DDR2/DDR3 SDRAM controllers.

To further speed up processing while keeping power dissipation down, the MPC8572E integrates powerful engines: a security engine that accelerates crypto operations in IPsec and SSL/TLS, a pattern-matching engine to handle regular expression matching, a deflate engine to manage file decompression and two table lookup units (TLU) that manage complex table searches and header inspections.

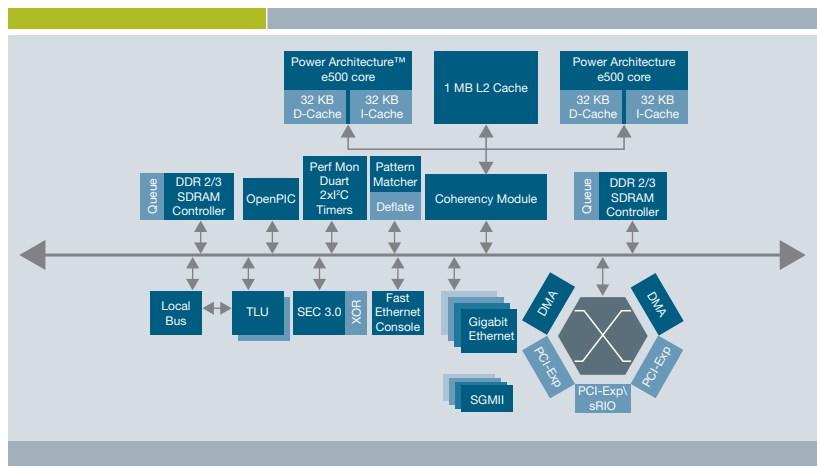
The MPC8572E offers a combination of network interfaces, including four integrated enhanced Triple-Speed Ethernet controllers (eTSEC). These controllers accelerate packet I/O by offloading checksum calculation. They also provide QoS support with eight Rx and eight Tx hardware queues to accelerate traffic management.

For high-speed connectivity to other devices, the MPC8572E supports PCI Express®, Serial RapidIO® and DMA interfaces.

All major processing and I/O elements are integrated into the MPC8572E with a highly optimized internal interconnect architecture to ensure high bandwidth, low latency and efficient pipeline operation, balancing processing performance with I/O system throughput.

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.

Figure 4: MPC8572E PowerQUICC® III Block Diagram



3.1 Pattern Matcher

The Pattern Matcher is the key contributor to the MPC8572E's ability to process packet content at high speed for application-aware networking and content security applications.

The Pattern Matcher is an integrated hardware block inside the MPC8572E with the following capabilities:

- High-performance, feature-rich hardware pattern matching of compressed and uncompressed data
 - Patterns expressed in Regular Expression (regex) with significant capabilities beyond that provided by the regex language
 - Stateful Rule—correlates multiple pattern matches and maintains state between matches
- Improvements over other pattern matching technologies:
 - No pattern “explosion” to support “wildcarding” or case-insensitivity
 - Fast compilation of pattern database
 - Fast incremental additions to pattern database
 - Live pattern database update
 - Patterns stored in main DDR DRAM, not SRAM or FCRAM
- On-chip hash tables for low system memory utilization, removing need for costly low-latency memory technologies
- Pattern matching across data “work units” (e.g. can match patterns split across TCP segments)

² Ptacek and Newsham, “Insertion, Evasion, and Denial of Service: Eluding Network Intrusion Detection”, 1998.

4 Designing IDS/IPS with the MPC8572E

4.1 IDS/IPS Processing

In order to understand what an ideal IDS/IPS platform looks like, let's examine the key operations performed in a typical IDS/IPS device.

IDS/IPS processing primarily consists of a detection phase followed by a response phase. The detection phase is concerned with the reception and scanning of data for malicious content, while the response phase is concerned with the steps to take after an attack has been detected.

The detection phase typically includes processing roughly in this order:

- Packets are received
- Received packets are classified into flows
- IP fragment reassembly is performed with careful scrutiny to detect IP fragment-based evasions/attacks² while reassembling
- TCP segment reassembly into a content byte stream is performed with careful scrutiny to detect TCP segment-based evasions/attacks² while reassembling
- Packet protocol anomalies are detected
- Patterns are matched on Layer 3 and 4 headers (against the header portion of the attack signatures)
- Application protocol anomalies are detected
- Stateful/context-based patterns are matched on the relevant portion of the content byte stream within and across packet boundaries (against the content portion of the relevant attack signatures)

Once an attack has been detected, an intrusion detection system must record the attack and generate an alert, if appropriate, while an intrusion prevention system must also block the attack:

- The packet/content is logged for suspicious flows
- Alerts are generated with correlation and filtering
- Alerts are sent to the management station
- For intrusion prevention, harmful content must be dropped while harmless content is packetized and forwarded to the destination

We observe that IDS/IPS datapath can be separated into

- packet processing
- content processing

Packet processing in the IDS/IPS looks similar, but is not quite the same as the relatively standard processing found in a host computer or bridge/router. One key difference is due to the fact that the NIDS/NIPS needs to “think like a hacker” to deal with evasion techniques used in fragmentation based attacks² for example. In fact, some vendors regard this as a highly proprietary part of their intellectual property. The packet processing capability in the platform needs to be flexible and addition to being fast in order to enable IDS/IPS vendors to implement their diversified, proprietary packet processing algorithms.


The most time-consuming part of application content processing – for that matter, all of NIDS/NIPS processing—is the matching of the appropriate part of network traffic against thousands of attack signatures in a stateful, context-sensitive manner.

² Ptacek and Newsham, “Insertion, Evasion, and Denial of Service: Eluding Network Intrusion Detection”, 1998.

4.2 IDS/IPS Operations on the MPC8572E

The dual-core MPC8572E can be used either in the Symmetric Multi-Processing (SMP) or the Asymmetric Multi-Processing (AMP) mode. The AMP is used in the rest of this white paper for illustrative purposes:

Figure 5

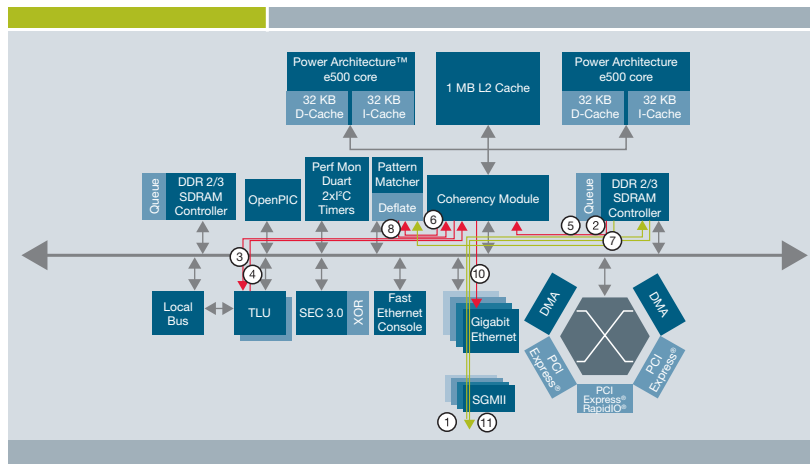


- Packet Processing performed on CPU Core1, with the TLU used to accelerate flow table lookup and eTSEC, network I/O
- Application Processing performed on CPU Core2, with the Pattern Matcher used to accelerate matching of packet content against attack signatures

In more detail, the IDS/IPS data path is shown in Figure 6:

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 lookup results
5. Core1 performs fragment & segment reassembly, checks for packet level anomalies and interrupts Core2
6. Core2 normalizes application data and instructs Pattern Matcher to scan content against Intrusion signatures
7. Pattern Matcher reads data from memory and scans for patterns
8. Pattern Matcher informs Core2 of (-ve) scan result
9. Core2 interrupts Core1
10. Core1 instructs eTSEC to transmit original packets
11. Data retrieved from memory and transmitted

Figure 6: MPC8572E



4.3 Performance Advantages

4.3.1 Packet I/O

The IDS/IPS—especially the IPS—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 device. The integrated eTSEC (Enhanced Triple Speed Ethernet Controller) is able to coalesce interrupts, thereby reducing the interrupt servicing overhead and improve performance.

In a networking device where the I/O rate is high, memory access speed in addition to the availability of CPU cycles can have a big impact on 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.

4.3.2 Packet Processing

A powerful e500 CPU core with clock speeds up to 1.5 GHz is dedicated to provide the CPU cycles (and flexibility) required for IDS/IPS packet-layer processing. Furthermore, the following operations are offloaded from the CPU core:

- IP and TCP checksum calculations to the eTSEC
- Flow table lookup to the TLU

Checksum calculations are required for every packet received. Offloading this calculation results in less software execution and higher performance.

As described earlier in Section 4.2, 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 Table Lookup Unit (TLU) provides hardware acceleration to this operation.

4.3.3 Application Content Processing

A powerful e500 core, working in conjunction with the hardware Pattern Matcher, enables high-speed analysis of the application protocols and their content.

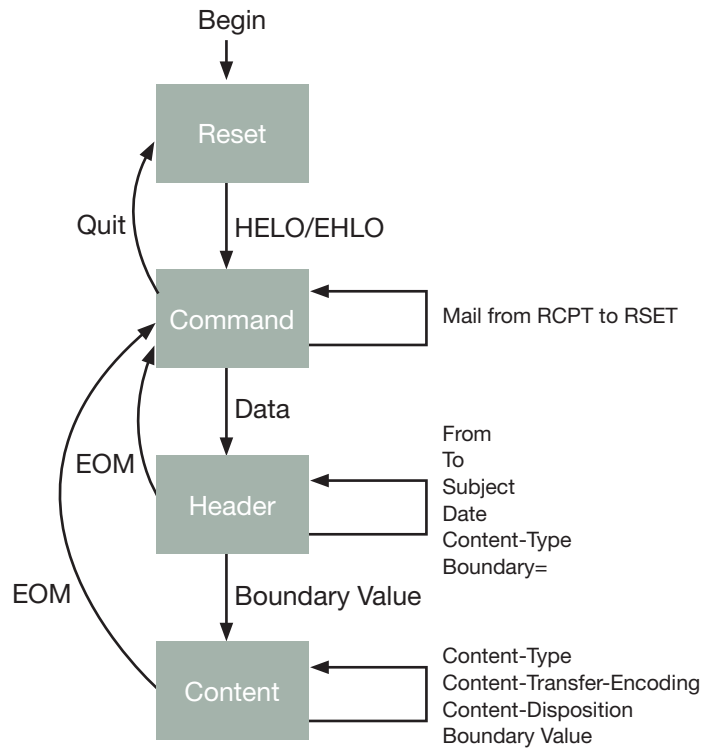
4.3.4 Stateful Pattern Matching

The most CPU-intensive operation in the IDS/IPS is the matching of appropriate portions of the network data against thousands of attack signatures in a stateful, context-sensitive manner.

Let us use an example to illustrate the implementation approaches of this CPU-intensive operation.

The simplified state transition diagram (see Figure 7) illustrates normal SMTP application protocol transitions, with each bubble representing a state and the text associated with an arrow representing an application message that would cause the transition from one particular state to another.

Figure 7



A good IDS/IPS would perform two distinct but related stateful pattern matching operations:

1. Track the state transition of the application protocol (and take appropriate action when an anomaly is detected)
2. Match packet content against signatures that are only appropriate for the particular state

There are three implementation approaches depending on the nature and capability of pattern matching in the platform:

1. A pure software implementation would perform both types of operations in software resulting in moderate performance
2. Most implementations with hardware pattern matching will perform application protocol tracking in software and matching of signatures in hardware resulting in improved performance
3. The MPC8572E's integrated Pattern Matcher, with its "stateful rule" and "regex" matching capability, can perform both types of operations resulting in highest performance

Incidentally, to ensure high performance, the interaction between the e500 core and the Pattern Matcher is very efficient—the two parties communicate efficiently via descriptors in L2 cache.

4.4 Accuracy (with Performance) Advantages

An IDS/IPS that provides false information about attacks is next to useless. Accuracy without performance, e.g. as in an IPS that throttles normal traffic, is also not acceptable.

The two powerful e500 CPU cores in the MPC8572E provide plenty of CPU cycles for the execution of software for the accurate detection of packet and content layer attacks. Furthermore, the built-in Pattern Matcher that offloads and accelerates pattern matching has a number of features and operational characteristics that are conducive to accuracy in intrusion detection:

- Regex
- Stateful Rule
- Set and subset
- Matching across packet boundaries
- Performance minimally dependent on number of signatures

4.4.1 Regex

The Regex Compiler associated with MPC8572E's Pattern Matcher supports a major subset of the PERL regular expression syntax, as well as capabilities beyond that provided by PERL. And the Pattern Matcher can match thousands of regexes in parallel at multi-Gbps speed.

Pattern matching based on regex is much more sophisticated than literal string match. This means that the signature designer has a much more powerful tool at his or her disposal to design sophisticated signatures to achieve high accuracy without worrying about performance.

4.4.2 Stateful Rule

As illustrated in Section 4.3.4, the Pattern Matcher's stateful rule capability can be used to track application protocols and create the context for stateful pattern matching to achieve high accuracy with performance.

4.4.3 Set and subset

Let's continue to use Figure 7 in Section 4.3.4 as an example. And let's say the connection is currently at the "Command" state. Only the applicable set/subset of the total signatures relevant to this state should be used to match the relevant portion of the packet content or else false positives may result. MPC8572E's Pattern Matcher supports sets and subsets.

Incidentally, for software pattern matching implementation, the performance usually drops significantly with the number of signatures configured. Using set/subset therefore increases accuracy and performance. For MPC8572E's Pattern Matcher, there is no significant dependency on the number of signatures configured.

4.4.4 Matching across packet boundaries

Application messages do not respect packet boundaries. As a result, an application layer attack signature – especially one carefully crafted by a knowledgeable hacker – can also span packet boundaries. To avoid false positives, matching across packet boundaries is required. This capability is supported in MPC8572E's Pattern Matcher.

4.4.5 Performance minimally dependent on number of signatures

As mentioned in Section 2, current Intrusion Systems often put IT managers in an awkward position to choose between accuracy or speed:

- Configuring relatively few signatures to achieve high performance at the risk of missing attacks that the complete set is capable of catching
- Configuring the complete set of signatures to detect all the known attacks but suffer from low performance as a result

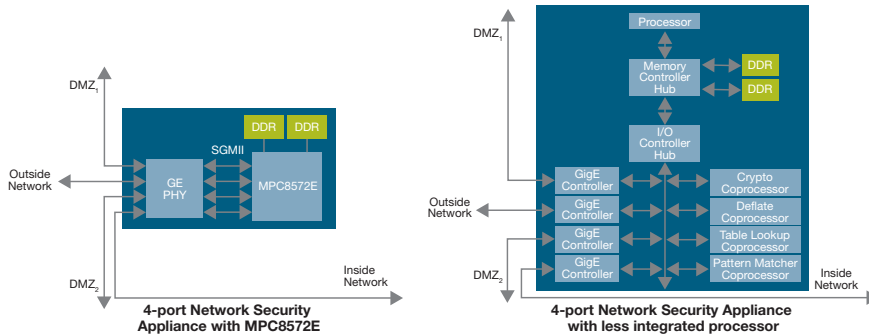
To solve this problem, the throughput performance of MPC8572E's integrated Pattern Matcher is minimally affected by the number of patterns configured.

4.5 Hardware Platform Design with the MPC8572E

In essence, the IPS is essentially a networking device that receives packets, processes the packet's header and the application content in the packet payload, and transmits the packet.

Figure 8 shows two simplified block diagrams of a 4-port network appliance, one implemented using the MPC8572E and the other, a less integrated processor.

Figure 8: 4-Port Networking Appliance Implementation



4.5.1 Cost Advantages

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.

There is no separate memory controller hub, I/O controller hub, Gigabit Ethernet controllers, table lookup coprocessor and pattern matcher coprocessor to complicate the design and add to the cost.

Specific to the Pattern Matcher, there is also no separate expensive low latency memory—the MPC8572E's built-in Pattern Matcher does not need it for high performance, unlike other pattern matching engines on the market.

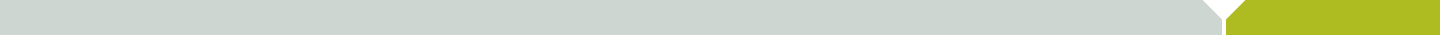
5 Summary

Freescale's MPC8572E, a PowerQUICC III processor optimized for Gbps network security and application-aware networking operations, is the perfect vehicle for OEMs to quickly deliver cost-effective Intrusion Detection/Prevention Systems with simultaneous high-performance and high-accuracy.

The MPC8572E's dual e500 cores provide CPU cycles and flexibility to execute the software to detect packet and application layer attacks. The CPU cores further orchestrate the integrated hardware blocks – eTSEC, TLU and Pattern Matcher in particular—to offload and accelerate CPU-intensive operations with low power dissipation.

The built-in Pattern Matcher, with its features and operational characteristics, is particularly conducive to providing high performance and accuracy simultaneously in intrusion detection and prevention:

- Regex for the creation of sophisticated attack signatures
- Stateful Rule to enable application protocol tracking and stateful pattern matching both in hardware
- Set and subset to enhance accuracy by comparing data against only signatures relevant to the state
- Matching across packet boundaries to increase accuracy by catching attacks that span packet boundaries
- Performance minimally dependent on number of signatures enables the IT manager to configure the complete set of signatures to catch all known attacks without worrying about degradation of performance



The MPC8572E processor—with all major processing and I/O elements included—enables very simple, elegant system design with low system cost and shortened design cycle. Contributing further to cost-effectiveness is the Pattern Matcher's use of DRAM instead of expensive low latency SRAM or FCRAM.

As a result, OEMs can count on using the MPC8572E to deliver highly competitive IDS/IPS products to the market place.

While this white paper focuses on describing how the MPC8572E can be used for IDS/IPS design, complementary white papers will show that the SoC can also be used for other security devices, including Firewall/VPN, anti-virus/anti-spam/content filter 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

Power Architecture Information:

www.freescale.com/powerarchitecture

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.