



The following excerpt is a complete chapter from The Linley Group's report, "[*A Guide to Security and Content Processors, Fourth Edition.*](#)"

This chapter is representative of the depth and quality that you will find in all of our reports. For a full list of current titles, please visit our [report page](#).
<http://www.linleygroup.com/reports.html>

For further information, contact:

The Linley Group
Phone: 408-281-1947
Fax: 650-745-1490
Email: cs@linleygroup.com
Website: www.linleygroup.com

11 Seaway

Company Background

Seaway Networks (Ottawa, Canada) is a fabless semiconductor startup developing content-processing silicon. Seaway was founded in 1Q01 by a team of engineers developing content-processing solutions at Nortel. Much of the management team comes from Nortel, including CEO Kit Fung and CTO David Lapp. In June 2004, Seaway closed a mezzanine funding round of about US\$8 million, bringing the total investment to around US\$23 million. Second-round investors included several participants in the first round, including JK&B Capital, VenGrowth, Venture Coaches, and Novacap. Seaway has been careful with expenses, maintaining its headcount at around 43 employees for the past year, partly by relying on contractors for portions of its software development.

Seaway's product is designed for deep packet processing as well as for application processing. The company's architecture, Streamwise, works on TCP flows, which Seaway calls data streams. Seaway sampled its first device, the SW5000, in May 2003. This device can terminate TCP, splice Layer 7 data streams, and assist with content processing. The company targets security equipment and data-center infrastructure, such as content switches, application firewalls, load balancers, SAN equipment, SSL accelerators, and intrusion-detection/prevention systems.

During 2004, Seaway introduced the NCA2000 family of PCI-X cards, targeting integrated security appliances. The new cards combine the SW5000 with processors for TCP acceleration, content processing, and cryptography in addition to substantial amounts of memory. Seaway now offers both boards and chips to its customers, but the company has not released new silicon since its initial product.

Key Features and Performance

Seaway calls its first product, the [SW5000](#), a Network Content Processor (NCP). Based on the [Streamwise](#) architecture, the SW5000 supports up to 5Gbps of throughput for Layer 4–7 processing. The SW5000 requires

external packet-control processors (PCP) to terminate Layer 4 protocols such as TCP, UDP, and ICMP. The company provides Layer 4 software for the Freescale MPC74xx processors,* which can be used as PCPs. By using a standard CPU for packet processing, Seaway and its customers can take advantage of standard development tools and the economies of scale of a major CPU vendor. Although TCP termination uses both the SW5000 and the MPC74xx, all packet processing is supervised by the PowerPC software, which can be quickly updated for protocol revisions or OEM customization. The TCP stack and content-processing firmware are maintained in the MPC74xx's on-chip caches, improving TCP and application-level performance.

Using a 500MHz PowerPC CPU such as the MPC7410, the SW5000 can terminate TCP at up to 2.5Gbps in each direction, enough for two Gigabit Ethernet (GbE) ports or one OC-48 port. With a 733MHz external processor, such as the MPC7457, the chip supports four Gigabit Ethernet ports or two OC-48 ports.

A system designer can use a second MPC74xx as the content-control processor (CCP), to manage stream processing. Alternatively, the system designer can omit the CCP and instead use the host processor for stream management. Because the host lacks access to some SW5000 features, Seaway recommends using a CCP for extensive processing; the company provides content-processing firmware for the MPC74xx.

The SW5000 assists the external host or content-control processor in performing content processing, such as examining, modifying, and replicating data streams. Examination includes identifying HTTP headers and cookies, filtering content, and even looking for viruses. Applications such as firewalls may take advantage of the chip's data-stream modification features. Examples of data-stream modification include encryption, network address translation (NAT), and cookie insertion. The SW5000 assists in these operations by providing the attached content processor with a stream of application-level data. Replication is useful for application-level multicasting, which is required in a web cache or streaming-media server.

The SW5000 has two network interfaces that can be configured either for POS-PHY Level 3 (PL3) or for a mode that connects to Hifn's IPsec or SSL accelerators (see [Chapter 8](#)) via a streaming data bus at up to 133MHz. Each PL3 interface can connect to either a single OC-48 framer or two GbE MACs. The interface supports two physical channels for Gigabit Ethernet but may require glue logic to connect to an external MAC.

The chip's three PowerPC processor interfaces connect to the optional CCP and one or two PCPs. The host processor must connect via a proprietary bus. Seaway provides FPGA-based bridge reference designs to convert the proprietary host bus to either PCI-X or HyperTransport (HT). The SW5000 uses external DDR memory to store data context and streams.

The SW5000 assists setup and teardown for up to 300,000 Layer 4 (TCP) connections per second. The chip can manage an impressive two million

*Complete coverage of Freescale's MPC74xx processors appears in our report *A Guide to High-Speed Embedded Processors*.

simultaneous connections, enough to support thousands of inactive connections with granularity for wireless subscribers; the chip also supports up to 64K virtual contexts for security or address domains.

The SW5000 can process up to 66 million ACL rules per second without PCP or CCP assistance. The device attains up to 600,000 Layer 7 switching decisions per second when a CCP is present; it can sustain 4.5Gbps on Layer 5–7 pattern-matching searches when the CCP uses the SW5000's hardware stream-searching assist.

The SW5000 is built in LSI Logic's 0.18-micron process and packaged in a 1,157-contact flip-chip BGA. It has a power dissipation of 8W (maximum). Seaway lists the device for \$645 in 1,000-unit quantities.

The NCA2000 content accelerator is a PCI-X card that targets integrated security appliances. The NCA2000 packs a large number of devices, including two Freescale MPC7457 processors, one each for packet processing and content processing; a Hifn 8154 security processor with up to 256MB of context memory; dual GbE interfaces (either copper or SFP optical); and up to 6GB of ECC DDR memory. Seaway offers four versions of the NCA2000 with different stuffing options—the crypto and content processors may each be omitted. Pricing ranges from \$4,000 for the version that excludes crypto and content acceleration to \$8,000 for the full-featured version.

Internal Architecture

Figure 11-1 is a block diagram of the SW5000, showing Layer 2–4 processing at the bottom left and right and upper-layer processing at the top. The chip consists of packet-transfer engines, a content-transfer engine, and a stream switch. The packet-transfer engines translate between streams and packets; the stream switch manages data movement inside the chip. The content-transfer engine moves stream data to the content-control processor, a host processor, or one or more coprocessors.

The chip has two packet-transfer engines (PTE) that combine with one or two external packet-control processors to terminate Layer 4 protocols. Under PCP supervision, each PTE performs integrity checks, validates IP headers, identifies TCP states, sends acknowledgments, and maintains segment order. The PTE also parses Layer 2, 3, and 4 headers and uses lookup tables for packet or stream classification. Layer 2 classification uses a table stored in the attached SDRAM to compare a combination of MAC source address, MAC destination address, and VLAN ID. The result is a “handle” or virtual interface that can be used by a security domain or by a Layer 3 domain.

By the time the PCP firmware is invoked, the PTE provides all relevant information regarding the packet, including the context, directly through memory-mapped registers. This feature reduces the PCP's packet-processing workload.

The PTE uses four tables for Layer 3–4 classification and transfers the results (except for discards) to the PCP firmware. Stored in external DDR memory, the tables are called source address reject, flow, listen, and access control list (ACL). An incoming packet is discarded if its source IP address matches a reject table entry. The flow table holds the five-tuple

(for TCP or UDP) or three-tuple (for other Layer 4 protocols) entries for existing streams; a matching packet is added to the existing flow. The flow entries in the listen table are set by the listen command over the sockets API. If no match exists in the flow table, the key is compared against the entries in the listen table. For a matching key, the SW5000 automatically creates a flow for the packet stream. The configurable ACL table provides 66Mpps (million searches per second) performance for up to 128-bit entries.

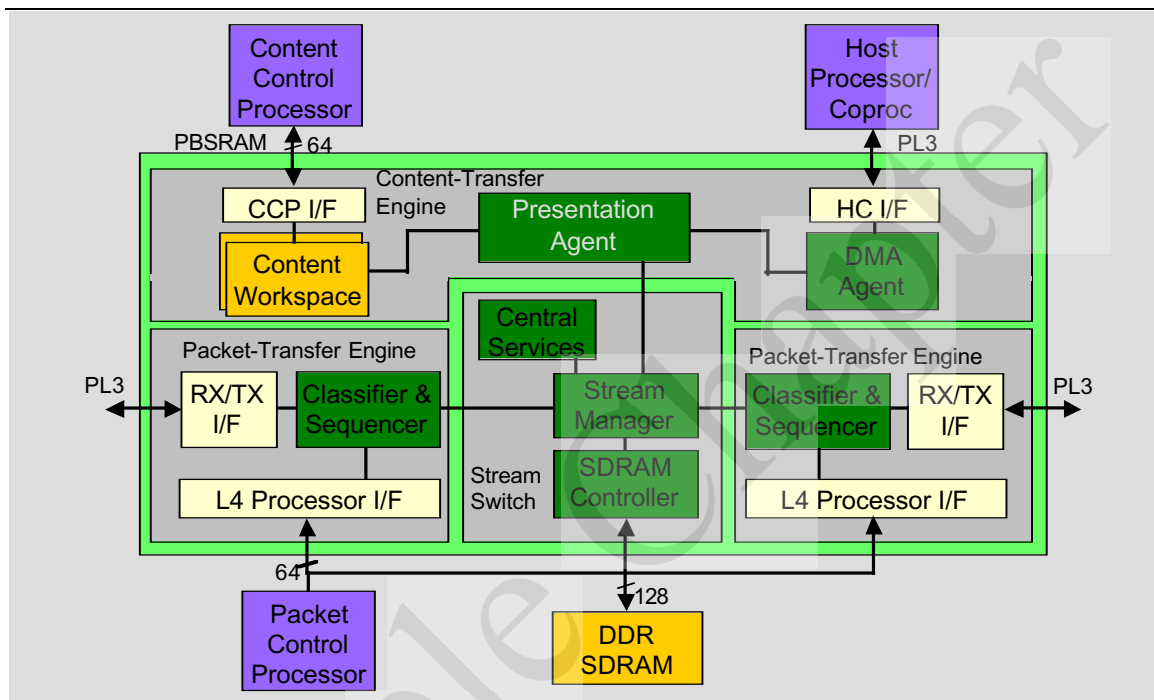


Figure 11-1. Seaway SW5000 internal architecture.

The PTE presents its lookup results and the first 224 bytes of the packet to the PCP. The PCP firmware determines the treatment appropriate for that flow. For example, it can decide to terminate TCP, as might be required for a new connection, or it can forward the packet payload after performing NAT, as might be required in a firewall.

The content-transfer engine (CTE) helps an external content processor examine, modify, and replicate streams. The SW5000 provides two types of assists: stream processing to accelerate data movement within the device and workspace management to accelerate content examination and modification. Examples of data-movement assists include adding data to the start or end of a stream; replicating a stream; removing or adding data to a stream; and moving DDR memory data into an existing stream.

The chip uses two dual-ported memories to present the stream to the external content processor. One memory is visible to the content processor; the other buffers data from the next stream. Because the content processor can switch between the two workspaces, it is fully utilized and unaffected by internal data movement. The SW5000 provides several

additional workspace-processing assists for the CCP only. CCP firmware invokes these features, including “search the stream for a pattern,” which is useful for filtering content or identifying viruses. Another example is “insert bytes,” which is useful for cookie insertion.

At the center of the SW5000 is the stream switch, which switches data among the chip’s function blocks as well as among the external interfaces. It can direct an incoming stream to an outgoing port, application processor, or coprocessor. It can also replicate the stream into multiple outgoing streams. For each Layer 4 session, the SW5000 creates an incoming and an outgoing stream, plus storage streams to move data objects (as large as 16MB) into external memory. In total, the stream switch provides a throughput of 20Gbps to support up to four million streams.

The stream switch can send a stream to multiple destinations or switch it multiple times. The latter case allows the same stream to be pipelined through multiple types of processing. For example, the content processor can examine a stream after an SSL coprocessor decrypts it.

Each stream-switch destination uses eight prioritized queues to buffer streams that require processing. A system designer can assign to the stream one of eight priorities, based on the DiffServ code point, 802.1p, or Layer 3 and 4 addresses. Once a stream meets its predefined eligibility criteria, it is scheduled into the appropriate priority queue of the destination. Examples of eligibility criteria include the amount of data available or a predefined event at the source of the stream. The destination of the stream can change the eligibility criteria, allowing the switch to redirect the stream. For example, a content processor could redirect a portion of the stream to an SSL coprocessor.

System Design

The SW5000 provides three 64-bit-wide processor interfaces that emulate pipelined burst SRAM and operate at 200MHz, providing 12Gbps of total peak bandwidth. These interfaces connect gluelessly to the external cache bus of Freescale’s MPC7410, 7455, or 7457 PowerPC processors, which are currently available at speeds ranging from 400MHz to 1.4GHz. Seaway recommends using a 733MHz MPC7457, which has a list price of about \$89 and power dissipation of 5W (typical); OEMs can choose a faster device for the CCP if additional performance is needed for customer-specific application code. Although IBM offers a series of software-compatible PowerPC processors, IBM’s processors do not offer an external cache bus and therefore cannot connect to the SW5000.

The SW5000 provides two equivalent ways to attach PCPs. The SW5000 can connect to separate PCPs, one per interface, or it can connect both interfaces to a single higher-performance PowerPC, using simple address-multiplexing logic to select between the two interfaces. Given the cost and power for an additional PowerPC processor and its associated memory devices, most customers will prefer to use a single processor of the appropriate speed unless they need maximum performance.

The host processor connects to a proprietary interface through an FPGA. Seaway provides FPGA code to bridge this interface to either PCI-X or HyperTransport.

The host processor can perform content processing. It has access to most of the SW5000's features, except for some content-workspace features. At 4.25Gbps full duplex, however, the host-interface bandwidth is less than that of the content-control processor interface. Seaway recommends using a CCP for applications requiring advanced content processing. The host (or coprocessor) interface supports scatter/gather for DMA to and from the host memory. The host interface can also be used in a PL3 mode or in a Hifn streaming mode, either of which allows direct connection with coprocessors, such as a security processor.

The SW5000 uses up to 4GB of DDR SDRAM memory to store data streams, lookup tables, and flow context. The amount of memory needed depends on the number of ports supported and the protocol to be terminated: applications using two Gigabit Ethernet ports are likely to need 512MB to 2GB; dual OC-48 applications are likely to need 1GB to 4GB.

Figure 11-2 shows the SW5000 in a multifunction security system with two Gigabit Ethernet ports. The system in this example might provide an application firewall and Layer 7 load balancing with SSL support. This example uses only one packet processor (owing to the slower wire speed) plus a content processor and an SSL coprocessor. The SW5000 supports security processors from all the major vendors and works gluelessly with security processors from Cavium, Corrent, and Hifn. For example, SafeNet's 1842 connects directly to the coprocessor PL3 port, as Figure 11-2 illustrates. A 2×GbE MAC, such as PMC-Sierra's PM3386, connects directly to the SW5000's packet-interface PL3 port.

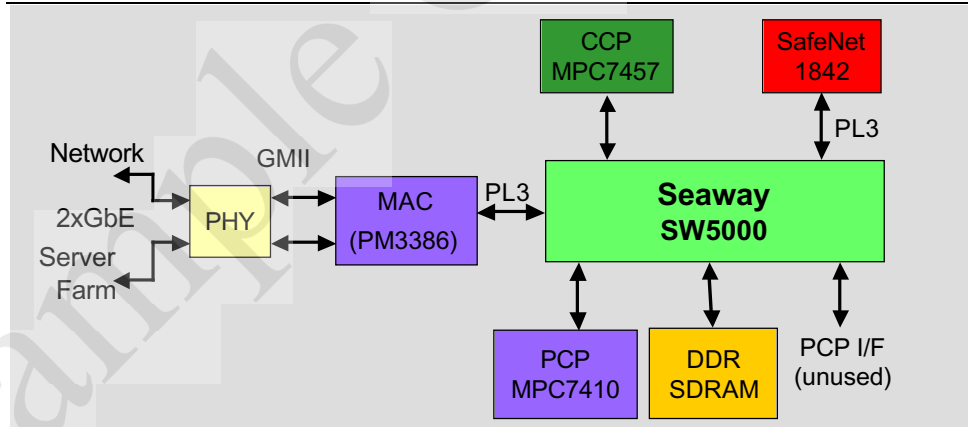


Figure 11-2. Seaway SW5000-based multifunction security system.

For an SSL transaction coming from the network, the SW5000 (assisted by the PCP) terminates the TCP connection and passes data streams to the CCP. After the CCP classifies the data stream, the SW5000 sends it to the SSL coprocessor for decryption. The decrypted stream then goes back to the CCP, which uses a load-balancing algorithm to switch the stream to an outgoing server-side port. The SW5000's packet-transfer engine converts the stream back to TCP packets for transmission to the network.

Figure 11-3 shows the SW5000 in a multifunction content blade. In this application, the chip receives data from interface cards in the system via the backplane. It connects to the switch-fabric interface through its two PL3 ports. For complex services, an OEM may add an external Layer 3 classification engine, via the PL3 port, to preprocess incoming traffic. The SW5000 terminates TCP and passes complete data streams to the host interface for processing.

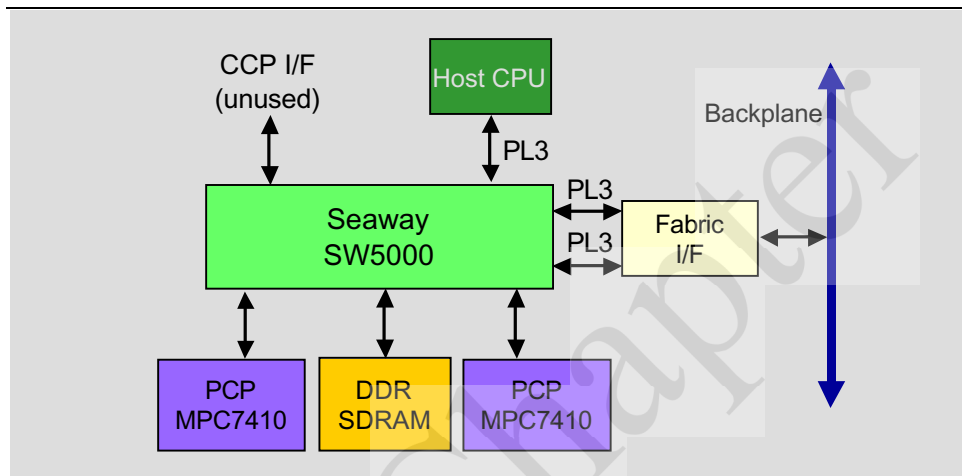


Figure 11-3. Seaway SW5000-based multifunction content blade.

If a security processor is required in this configuration, it must share the host interface, because both PL3 ports connect to the fabric. This sharing is accomplished using an FPGA; depending on the amount of SSL processing required, the FPGA can connect to more than one SSL coprocessor. Seaway provides reference code for various FPGA bridge designs: the SWi202, for instance, provides a PL3-to-PCI bridge. Other reference designs include a PL3-to-HT bridge and a dual PCI-X bridge. The SW5000 can keep track of multiple coprocessors and send data to the appropriate one for encryption or decryption, or for OEM-specific processing. After the data is processed, the SW5000 transmits the stream via the switch and the PTE.

The NCA2000, shown in Figure 11-4, uses a full-size PCI form factor and dissipates, by The Linley Group's estimates, more than 45W. Because this number exceeds the maximum allowed PCI-X power dissipation, the NCA2000 uses an auxiliary power cable to tap into a 12V power source. To address thermal concerns, Seaway has outfitted the NCA2000 with a custom-designed heat sink; still, cooling is an important consideration for system integrators. The NCA2000 requires adequate airflow, shutting itself down if heat-sink temperature monitors indicate that the limit has been exceeded. Seaway has qualified a number of popular appliance chassis for the NCA2000 card. Seaway also assists customers with qualifying other chassis.

The NCA2000 PCI-X Card Development Platform is a self-contained system for network security and content-processing application development. The development platform houses the NCA2000 in a standard 1U

rack-mount server chassis that has been prequalified for mechanical, thermal, power, and functional compatibilities. The platform is preloaded with the necessary system software as well as the tools required for application-software development and testing. Seaway provides drivers and firmware in source and object form.

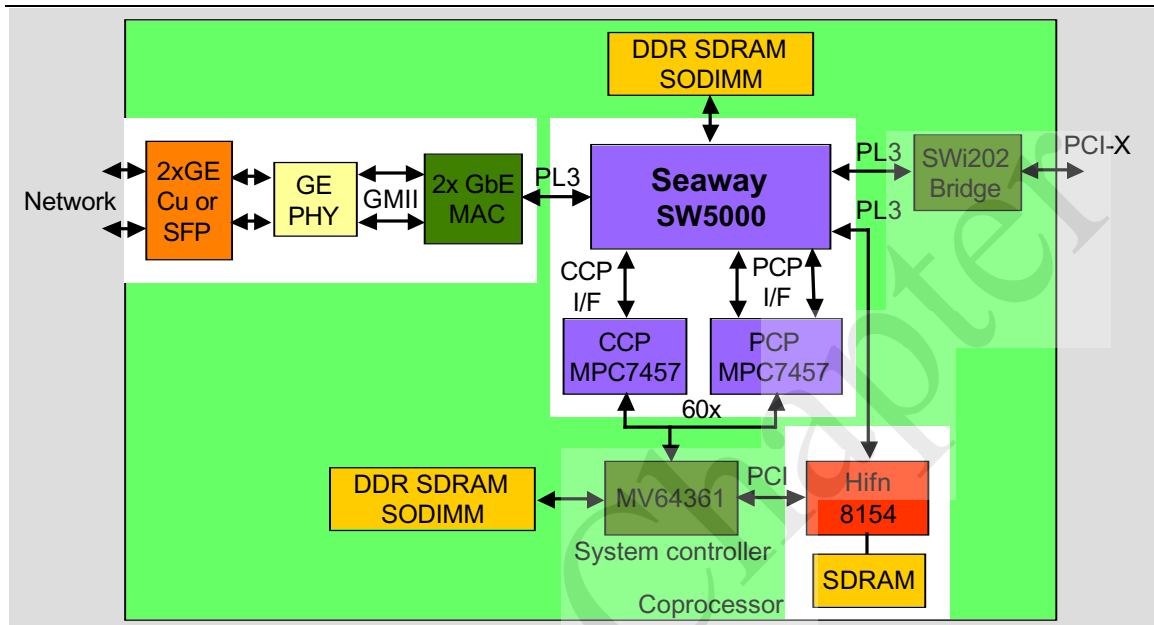


Figure 11-4. Seaway NCA2000 network-content accelerator card.

Seaway provides a hardware-development platform, priced at \$21,495, consisting of a 6U CompactPCI chassis with one SW5000 board and various plug-in options, such as GbE interface cards. The company provides a GNU PowerPC cross-compiler with related tool chain and software-development kit. This environment is also supported on Linux and Solaris and can be supported on Windows using Cygwin. The programmer is free to choose common open-source development tools, such as Vim, Emacs, XEmacs, or DDD, or an integrated development environment.

Seaway provides Layer 4 libraries and a sockets API for OEM software as well as reference applications targeting security. The company supplies the firmware as source or object code; source code allows customers to modify and add enhancements to the code, but Seaway charges a fee in this case because of additional support costs.

Product Roadmap

Seaway has not yet released information on future products; however, the company has indicated that future versions of the content accelerator board will include hardware support for wire-speed regular-expression evaluation. Seaway's planned regular-expression hardware is not based on deterministic finite automata (DFA). Instead, the company's undisclosed algorithm results in substantial memory savings for a regular-

expression subset that is sufficiently expressive for IDS/IPS applications. A single instance of Seaway's regular-expression processing block provides 4Gbps throughput at 300MHz.

Seaway plans to release a version of the NCA2000 that uses FPGAs to implement the regular-expression processing block. We expect this board to sample during 1H05.

Although the NCA2000 is an excellent platform to validate architectural approaches and facilitate application-porting efforts, its high cost and power dissipation will limit adoption. A natural step for Seaway is to pursue an integration strategy, combining the board-level elements into a single SoC. Seaway's next-generation silicon will likely take advantage of 90nm fabrication technology to integrate the packet- and content-processing general-purpose CPUs, cryptography, regular-expression engines, and network interfaces (MACs). Seaway's existing customers will readily adopt the new content processor to reduce appliance materials cost. Seaway has not provided a target sample date for the new device, but we do not expect the device to sample before late 2006.

Seaway's unique partitioning of separate processors for lower- and upper-layer processing makes the NCA2000 suited for applications other than integrated security. For instance, session border controllers for VoIP require Layer 4 hardware that can handle a large number of Real-Time Protocol (RTP) voice connections as well as a Layer 7 engine that can handle the Session Initiation Protocol (SIP) requirements of session setup and teardown. These requirements map conveniently to the PCP and CCP, respectively. The NCA2000 can handle several hundred thousand sessions compared with Pentium-based hardware that peaks at around 20,000 sessions. Seaway's pursuit of these other markets is likely to be constrained by internal development and support resources rather than by technology.

Conclusions

Seaway's Streamwise architecture takes a unique approach to content processing, providing a hardware data path that attaches processing tasks for lower- and upper-layer functions to separate general-purpose processors. For Layer 4 and below, Seaway's fast-path hardware includes some NPU-like features (hardware classification and ACL filtering, for instance) but differs in using a general-purpose processor to control TCP termination. Streamwise's biggest departure from traditional NPU architectures is its hardware support for streams.

Seaway's key contribution is creating the "stream" abstraction as well as an efficient means of routing these streams through any number of sequential software- or hardware-processing stages. The conversion between packets, which may arrive out of order and may contain only part of application-layer payloads, and streams—composed of ordered application-layer payloads and held in memory—results in more-efficient content processing. For example, the SW5000 can launch content-processing functions only when sufficient application-layer data exist, instead of every time a new packet arrives. Seaway's implementation of streams makes the SW5000 one of the most flexible content-processing

architectures on the market, at the cost of adding one or more general-purpose processors.

Despite its daunting price and high power dissipation, the NCA2000 is the only off-the-shelf card to provide access-control-list (ACL) filtering, full TCP offload, and C-programmable content inspection at 1Gbps full-duplex throughput. Security vendors that offer a Pentium-based appliance with a throughput of about 100Mbps can upgrade performance to 1Gbps by porting their application to the NCA2000 board. This \$8,000 board enables the vendor to charge about \$20,000 more for the appliance, reflecting the market premium for delivering gigabit throughput on content applications.

The NCA2000's use of heat sinks and memory modules presents difficulties in making the card tamperproof. Until Seaway provides a more integrated solution, the NCA2000 will not be selected for ultra-secure applications that mandate FIPS 140 Level 2 or Level 3 certification.

Seaway's partitioning of the problem between general-purpose processors and hardware is favorable to software developers, since many key content-processing routines can remain in a single high-level-language form. Security-software vendors benefit by developing a single program source for both accelerated and nonaccelerated versions. In particular, Seaway's use of a general-purpose PowerPC processor to control TCP termination provides security vendors with the ability to continue to evolve lower-layer code to counter developing threats; this adaptability would not be provided by a hard-wired TCP termination engine.

For deep content inspection, Seaway will need to add specific hardware acceleration to remain competitive with other vendors. Seaway plans to offer regular-expression processing hardware, initially at the board level. This new hardware will dilute the porting benefit of using general-purpose processors for content-inspection applications; however, unlike other vendors (such as Cavium) that have already committed silicon to an approach, Seaway will be at a cost disadvantage with an FPGA solution. Seaway claims its approach does not have the extravagant memory requirements of DFA-based approaches, thereby simplifying SoC integration and reducing system cost.

Seaway's use of general-purpose processors is a significant obstacle to integration. An off-the-shelf synthesized processor core will fall short of the performance of the standalone processors used in Seaway's boards. One choice would be to contract with IBM as the foundry for the next-generation chip and take advantage of IBM's PowerPC 440 CPU core. IBM is an expensive foundry, however.

Cavium's sampling of Oction, expected in 1H05, will increase the pressure on Seaway to offer an integrated solution; furthermore, Oction is based on custom-designed MIPS cores. Oction will also enable board-level competition for Seaway by the end of 2005.

Vendors of content-inspection cards compete for the attention of security-software vendors such as Symantec. Seaway, as a platform vendor, must provide compelling reasons to induce security-software vendors to port to its platform. The ease of porting to Seaway's hardware works to Seaway's advantage, but without additional customers and funding, Seaway is unlikely to attract attention from major software vendors.

Seaway's next phase of silicon integration is a challenging prospect for the startup; a processor company might be in a better position to take that step. Accordingly, an acquisition by a processor company is one possible exit scenario for Seaway. The acquiring company would gain the strategic value of Seaway's technology, enabling the acquirer to make a credible entry into the content-processing space. Because Seaway's approach is counter to the conventional NPU approach, the acquiring company might be one that lacks NPU technology but wishes to enter the networking market, or a company, such as Freescale, that has abandoned its NPU development. Absent such suitors, Seaway will have to remain a subsystem vendor and hope to generate enough revenue to fund its next development effort.