

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
12 July 2007 (12.07.2007)

PCT

(10) International Publication Number
WO 2007/079095 A2

(51) International Patent Classification:
G06F 15/16 (2006.01)

(21) International Application Number:
PCT/US2006/049317

(22) International Filing Date:
26 December 2006 (26.12.2006)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
11/323,165 30 December 2005 (30.12.2005) US

(71) Applicant and

(72) Inventor: **PANDYA, Ashish, A.** [US/US]; 4318 Lafayette Drive, El Dorado Hills, CA 95762 (US).

(74) Agent: **LEAL, Peter, R.**; DLA PIPER US LLP, ATTN: PATENT DEPARTMENT, 2000 University Avenue, East Palo Alto, CA 94303-2248 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM,

AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LV, LY, MA, MD, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.


(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished upon receipt of that report

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: **RUNTIME ADAPTABLE SEARCH PROCESSOR**



(57) Abstract: A runtime adaptable search processor is disclosed. The search processor provides high speed content search capability to meet the performance need of network line rates growing to 1Gbps, 10Gbps and higher, the search processor provides a unique combination of NFA and DFA based search engines that can process incoming data in parallel to perform the search against the specific rules programmed in the search engines. The processor architecture also provides capabilities to transport and process Internet Protocol (IP) packets from Layer 2 through transport protocol layer and may also provide packet inspection through Layer 7. Further, a runtime adaptable processor is coupled to the protocol processing hardware and may be dynamically adapted to perform hardware tasks as per the needs of the network traffic being sent or received

and/or the policies programmed or services or applications being supported. A set of engines may perform pass-through packet classification, policy processing and/or security processing enabling packet streaming through the architecture at nearly the full line rate. A high performance content search and rules processing security processor is disclosed which may be used for application layer and network layer security, scheduler schedules packets to packet processors for processing. An internal memory or local session database cache stores a session information database for a certain number of active sessions. The session information that is not in the internal memory is stored and retrieved to/from an additional memory. An application running on an initiator or target can in certain instantiations register a region of memory, which is made available to its peer(s) for access directly without substantial host intervention through RDMA data transfer. A security system is also disclosed that enables a new way of implementing security capabilities inside enterprise networks in a distributed manner using a protocol processing hardware with appropriate security features.

WO 2007/079095 A2

RUNTIME ADAPTABLE SEARCH PROCESSOR

RELATED APPLICATIONS

Priority is claimed to U.S. Patent Application Serial No. 11/323,165, filed on December 30, 2005 entitled Runtime Adaptable Search Processor.

5 BACKGROUND OF THE INVENTION

This invention relates generally to content search, storage and networking semiconductors and in particular to high performance content search, network storage and security processors that can be used within networking, storage, security, bioinformatics, chipsets, servers, search engines and the like.

- 10 Many modern applications depend on fast information search and retrieval. With the advent of the world-wide-web and the phenomenal growth in its usage, content search has become a critical capability. A large number of servers get deployed in web search applications due to the performance limitations of the state of the art microprocessors for regular expression driven search.
- 15 There have been significant research and development resources devoted to the topic of searching of lexical information or patterns in strings. Regular expressions have been used extensively since the mid 1950s to describe the patterns in strings for content search, lexical analysis, information retrieval systems and the like. Regular expressions were first studied by S.C. Kleene in mid-1950s to describe the events of nervous activity. It is well understood
- 20 in the industry that regular expression (RE) can also be represented using finite state automata (FSA). Non-deterministic finite state automaton (NFA) and deterministic finite state automaton (DFA) are two types of FSAs that have been used extensively over the history of computing. Rabin and Scott were the first to show the equivalence of DFA and NFA as far as their ability to recognize languages in 1959. In general a significant body of
- 25 research exists on regular expressions. Theory of regular expressions can be found in "Introduction to Automata Theory, Languages and Computation" by Hopcroft and Ullman and a significant discussion of the topics can also be found in book "Compilers: Principles, Techniques and Tools" by Aho, Sethi and Ullman.

- Internet protocol (IP) is the most prevalent networking protocol deployed across various
- 30 networks like local area networks (LANs), metro area networks (MANs) and wide area

networks (WANs). Storage area networks (SANs) are predominantly based on Fibre Channel (FC) technology. There is a need to create IP based storage networks.

When transporting block storage traffic on IP designed to transport data streams, the data streams are transported using Transmission Control Protocol (TCP) that is layered to run on top of IP. TCP/IP is a reliable connection/session oriented protocol implemented in software within the operating systems. TCP/IP software stack is very slow to handle the high line rates that will be deployed in future. Currently, a 1 GHz processor based server running TCP/IP stack, with a 1Gbps network connection, would use 50-70% or more of the processor cycles, leaving minimal cycles available for the processor to allocate to the applications that run on the server. This overhead is not tolerable when transporting storage data over TCP/IP as well as for high performance IP networks. Hence, new hardware solutions would accelerate the TCP/IP stack to carry storage and network data traffic and be competitive to FC based solutions. In addition to the TCP protocol, other protocols such as SCTP and UDP protocols can be used, as well as other protocols appropriate for transporting data streams.

Computers are increasingly networked within enterprises and around the world. These networked computers are changing the paradigm of information management and security. Vast amounts of information, including highly confidential, personal and sensitive information is now being generated, accessed and stored over the network, which information needs to be protected from unauthorized access. Further, there is a continuous onslaught of spam, viruses, and other inappropriate content on the users through email, web access, instant messaging, web download and other means, resulting in significant loss of productivity and resources.

Enterprise and service provider networks are rapidly evolving from 10/100Mbps line rates to 1Gbps, 10Gbps and higher line rates. Traditional model of perimeter security to protect information systems pose many issues due to the blurring boundary of an organization's perimeter. Today as employees, contractors, remote users, partners and customers require access to enterprise networks from outside, a perimeter security model is inadequate. This usage model poses serious security vulnerabilities to critical information and computing resources for these organizations. Thus the traditional model of perimeter security has to be bolstered with security at the core of the network. Further, the convergence of new sources of threats and high line rate networks is making software based perimeter security to stop the external and internal attacks inadequate. There is a clear need for enabling security processing in hardware inside core or end systems beside a perimeter firewall as one of the prominent means of security to thwart ever increasing security breaches and attacks.

FBI and other leading research institutions have reported in recent years that over 70% of intrusions in organizations have been internal. Hence a perimeter defense relying on protecting an organization from external attacks is not sufficient as discussed above.

Organizations are also required to screen outbound traffic to prevent accidental or malicious disclosure of proprietary and confidential information as well as to prevent its network resources from being used to proliferate spam, viruses, worms and other malware. There is a clear need to inspect the data payloads of the network traffic to protect and secure an organization's network for inbound and outbound security.

Data transported using TCP/IP or other protocols is processed at the source, the destination or intermediate systems in the network or a combination thereof to provide data security or other services like secure sockets layer (SSL) for socket layer security, Transport layer security, encryption/decryption, RDMA, RDMA security, application layer security, virtualization or higher application layer processing, which may further involve application level protocol processing (for example, protocol processing for HTTP, HTTPS, XML, SGML, Secure XML, other XML derivatives, Telnet, FTP, IP Storage, NFS, CIFS, DAFS, and the like). Many of these processing tasks put a significant burden on the host processor that can have a direct impact on the performance of applications and the hardware system. Hence, some of these tasks need to be accelerated using dedicated hardware for example SSL, or TLS acceleration. As the usage of XML increases for web applications, it is expected to put a significant performance burden on the host processor and would also benefit significantly from hardware acceleration. Detection of spam, viruses and other inappropriate content require deep packet inspection and analysis. Such tasks can put huge processing burden on the host processor and can substantially lower network line rate. Hence, deep packet content search and analysis hardware is also required.

Internet has become an essential tool for doing business at small to large organizations. HTML based static web is being transformed into a dynamic environment over last several years with deployment of XML based services. XML is becoming the lingua-franca of the web and its usage is expected to increase substantially. XML is a descriptive language that offers many advantages by making the documents self-describing for automated processing but is also known to cause huge performance overhead for best of class server processors. Decisions can be made by processing the intelligence embedded in XML documents to enable business to business transactions as well as other information exchange. However, due to the performance overload on the best of class server processors from analyzing XML documents, they cannot be used in systems that require network line rate XML processing to provide intelligent networking. There is a clear need for acceleration solutions for XML

document parsing and content inspection at network line rates which are approaching 1Gbps and 10Gbps, to realize the benefits of a dynamic web based on XML services.

- Regular expressions can be used to represent the content search strings for a variety of applications like those discussed above. A set of regular expressions can then form a rule set for searching for a specific application and can be applied to any document or stream of data for examination of the same. Regular expressions are used in describing anti-spam rules, anti-virus rules, XML document search constructs and the like. These expressions get converted into NFAs or DFAs for evaluation on a general purpose processor. However, significant performance and storage limitations arise for each type of the representation. For example an N character regular expression can take up to the order of 2^N memory for the states of a DFA, while the same for an NFA is in the order of N. On the other hand the performance for the DFA evaluation for an M byte input data stream is in the order of M memory accesses and the order of (N * M) processor cycles for the NFA representation on modern microprocessors.
- When the number of regular expressions increases, the impact on the performance deteriorates as well. For example, in an application like anti-spam, there may be hundreds of regular expression rules. These regular expressions can be evaluated on the server processors using individual NFAs or DFAs. It may also be possible to create a composite DFA to represent the rules. Assuming that there are X REs for an application, then a DFA based representation of each individual RE would result up to the order of $(X * 2^N)$ states however the evaluation time would grow up to the order of $(X * N)$ memory cycles. Generally, due to the potential expansion in the number of states for a DFA they would need to be stored in off chip memories. Using a typical access time latency of main memory systems of 100ns, it would require about $(X * 100ns * N * M)$ time to process an X RE DFA with N states over an M byte data stream. This can result in tens of Mbps performance for modest size of X, N & M. Such performance is obviously significantly below the needs of today's network line rates of 1Gbps to 10Gbps. On the other hand, if a composite DFA is created, it can result in an upper bound of storage in the order of $2^{N \times X}$ which may not be within physical limits of memory size for typical commercial computing systems even for a few hundred REs. Thus the upper bound in memory expansion for DFAs can be a significant issue. Then on the other hand NFAs are non-deterministic in nature and can result in multiple state transitions that can happen simultaneously. NFAs can only be processed on a state of the art microprocessor in a scalar fashion, resulting in multiple executions of the NFA for each of the enabled paths. X REs with N characters on average can be represented in the upper bound of $(X \times N)$ states as NFAs. However, each NFA would

require M iterations for an M-byte stream, causing an upper bound of $(X \cdot N \cdot M \cdot \text{processor cycles per loop})$. Assuming the number of processing cycles are in the order of 10 cycles, then for a best of class processor at 4 GHz, the processing time can be around $(X \cdot N \cdot M \cdot 2.5\text{ns})$, which for a nominal N of 8 and X in tens can result in below 100 Mbps performance. There is a clear need to create high performance regular expression based content search processors which can provide the performance in line with the network rates which are going to 1Gbps and 10Gbps.

The methods for converting a regular expression to NFA and DFA are well known. The resulting automata are able to distinguish whether a string belongs to the language defined by the regular expression however it is not very efficient to figure out if a specific sub-expression of a regular expression is in a matching string or the extent of the string. Tagged NFAs enable such queries to be conducted efficiently without having to scan the matching string again. For a discussion on Tagged NFA please refer to the paper "NFAs with Tagged Transitions, their Conversion to Deterministic Automata and Application to Regular Expressions", by Ville Laurikari, Helsinki University of Technology, Finland.

U.S. Patent Applications, 20040059443 and 20050012521, describe a method and apparatus for efficient implementation and evaluation of state machines and programmable finite state automata. These applications show an apparatus that is used to evaluate regular expressions using an array of NFAs to create high performance processing of regular expressions. The application recognizes the upper bound in the storage issues for DFAs as a reason to implement regular expressions using NFAs. However, the applications fails to recognize that even though the DFA worst case storage requirement is substantially higher compared to NFAs many DFAs have less storage needs than NFAs. DFAs for many regular expressions can result in lower number of states compared to an NFA. For example in an anti-spam application, based on the open source tool SpamAssassin, a large number of the regular expression rules result in DFAs which are smaller than NFAs. Hence, it is important not to ignore DFA implementation based only on the worst case scenario. These patent applications also create NFA engines that process a single RE per NFA block. Thus if a RE uses fewer states than the minimum states of the NFA block, there is no provision to be able to use multiple REs simultaneously in the same block. In my invention, I describe a content search processor which uses an array of runtime adaptable search engines, where the search engines may be runtime adaptable DFA search engines or runtime adaptable NFA search engines or a combination thereof to evaluate regular expressions. Content search engine of my search processor also provides flexibility of using multiple REs per NFA or DFA engine. My invention also provides capabilities to support Tagged NFA implementations

which are not supported or discussed in these applications. Further, these applications do not address the need of dynamically configuring the hardware or the rules being applied based on the transported data being sent to or received from a network. The processors of my invention can be dynamically adapted to apply hardware based rule sets dependent on the transported data which is not described in the above applications. Further, my invention shows that certain DFAs can be more hardware resource efficient to implement compared to NFAs and can enable today's state of the art FPGAs to implement a large number of regular expressions without having to devote large investments in creating application specific integrated circuits using advanced process technologies. This is also specifically discussed as not feasible to do in the above applications. My invention also shows content search acceleration can be used to improve application acceleration through content search application programmer interface (API) and the search processor of this invention.

Hardware acceleration for each type of network data payload can be expensive when a specialized accelerator is deployed for each individual type of network data. There is a clear need for a processor architecture that can adapt itself to the needs of the network data providing the necessary acceleration and thereby reduce the impact on the host performance. This patent describes such a novel architecture which adapts itself to needs of the network data. The processor of this patent can be reused and adapted for differing needs of the different types of the payload and still offer the benefits of hardware acceleration. This can have a significant reduction in the cost of the acceleration solutions deployment compared to dedicated application-specific accelerators.

Dynamically reconfigurable computing has been an area that has received significant research and development interest to address the need of reconfiguring hardware resources to suit application needs. The primary focus of the research has been towards creating general purpose microprocessor alternatives that can be adapted with new instruction execution resources to suit application needs.

Field programmable gate arrays (FPGA) have evolved from simple AND-OR logic blocks to more complex elements that provide a large number of programmable logic blocks and programmable routing resources to connect these together or to Input/Output blocks. U.S. patent No. 5,600,845 describes an integrated circuit computing device comprising a dynamically configurable FPGA. The gate array is configured to create a RISC processor with a configurable instruction execution unit. This dynamic re-configurability allows the dynamically reconfigurable instruction execution unit to be changed to implement operations in hardware which may be time consuming to run in software. Such an arrangement

requires a preconfigured instruction set to execute the incoming instruction and if an instruction is not present it has to be treated as an exception which then has a significant processing overhead. The invention in U.S. patent No. 5,600,845 addresses the limitation of general purpose microprocessors but does not address the need of dynamically configuring the hardware based on the transported data being sent to or received from a network.

U.S. Patent Application number 20030097546 describes a reconfigurable processor which receives an instruction stream that is inspected by a instruction test module to decide if the instruction is supported by existing non reconfigurable hardware or the reconfigurable hardware configured by a software routine and executes the instruction stream based on the test result. If the instruction is not supported then the processor decides a course of action to be taken including executing the instruction stream in software. The patent application number 20030097546 also does not address the need of dynamically configuring the hardware based on the transported data being sent to or received from a network.

U.S. Patent Application number 20040019765 describes a pipelined reconfigurable dynamic instruction set processor. In that application, dynamically reconfigurable pipeline stages under control of a microcontroller are described. This is yet another dynamically reconfigurable processor that can adapt its pipeline stages and their interconnections based on the instructions being processed as an alternative to general purpose microprocessors.

The field of reconfigurable computing has been ripe with research towards creating dynamically reconfigurable logic devices either as FPGAs or reconfigurable processors as described above as primarily addressing the limitations of general purpose processors by adding reconfigurable execution units or reconfigurable coprocessors. For example, "Reconfigurable FPGA processor", diploma thesis paper by Andreas Romer from Swiss Federal Institute of Technology, targets the need of creating an ASIC-like performance and area, but general purpose processor level flexibility, by dynamically creating execution functional units in a reconfigurable part of a reconfigurable FPGA like Xilinx Virtex and XC6200 devices. Similarly, the paper by J.R. Hauser and J Wawrzynek entitled *Garp: A MIPS Processor With a Reconfigurable Coprocessor* published in Proceedings of the IEEE Symposium on FPGAs for Custom Computing Machines (FCCM '97), targets the need for creating custom co-processing support to a MIPS processor addressing the limitations of the general purpose processing capabilities of the MIPS processor.

Published research or patent applications have not addressed the need of dynamically configuring the hardware based on transported data as well as actions to be taken and

applications/services to be deployed for that specific data being sent to or received from a network. This patent describes a novel architecture which adapts itself to the needs of the network data and is run-time adaptable to perform time consuming security policy operations or application/services or other data processing needs of the transported data and defined policies of the system incorporating this invention. The architecture also comprises a deep packet inspection engine that may be used for detecting spam, viruses, digital rights management information, instant message inspection, URL matching, application detection, malicious content, and other content and applying specific rules which may enable anti-spam, anti-virus and the like capabilities.

10 SUMMARY OF THE INVENTION

I describe a high performance run time adaptable search processor using hardware acceleration for regular expressions. The regular expressions are converted into equivalent DFAs and NFAs and then the most cost effective solution is chosen. The content search processor comprises of a set of DFA processing engines and NFA processing engines. The converted REs are then mapped either to an appropriate NFA or DFA engine. The processor may also include a set of composite DFA processing engines which can be used to absorb the growth in the number of rules beyond the number supported by the array of FSA engines. Thus new system hardware may not be necessary until the composite DFA results in significant memory usage which is beyond that available with the memory associated with the content search processor.

In many content search applications like security, there is a need to constantly update the rules or the signatures being used to detect malicious traffic. In such applications it is critical that a solution be adaptable to keep up with the constantly evolving nature of the security threat. In an always connected type of usage models, it is extremely important to have the latest security threat mitigation rules updated in the security system on a frequent basis. When a composite DFA type architecture is used, compiling and releasing any new security rules or policy can consume a large amount of time, where the updates may not be timely to avoid the impact of the security threat. In such environments the release of new rule base may take up to 8 to 24 hours, which is quite delayed response to constantly evolving threat. In the processor of this invention, that issue is addressed since the release of new rules is a matter of converting those rules into NFAs and DFAs and updating only these very small rules into the content search processor. Thus the response to new threats can be immediate and would not require huge delays which occur from integration of the new rules in the composite rule base and converting those into composite DFAs.

There are several instances of REs which include only a few states. For example if the content search includes looking for *.exe or *.com or *.html or the like, the NFA or DFAs for these REs include a small number of states. Thus if all DFA or NFA engines support say 16 states, then it may be possible to include multiple rules per engine. This invention enables
5 the maximum utilization of the FSA engines by allowing multiple rules per FSA engine. The engines also provide FSA extension logic to chain the base engines together to create super blocks that can handle larger FSAs.

Berry and Sethi in their paper "From Regular Expressions to Deterministic Automata" Published in Theoretical Computer Science in 1986, showed that regular expressions can be
10 represented by NFAs such that a given state in the state machine is entered by one symbol, unlike the Thompson NFA. Further, the Berry-Sethi NFAs are ϵ -free. A 'V' term RE can be represented using 'V+1' states NFA using Berry-Sethi like NFA realization method. The duality of Berry-Sethi method also exists where all transitions that lead the machine out of a state are dependent on the same symbol. This is shown in the paper "A Taxonomy of finite
15 automata construction algorithms" by Bruce Watson published in 1994 in section 4.3. I show a method of creating NFA search engine architecture leveraging the principles of Berry-Sethi's NFA realization and the dual of their construct. The NFA search engine is programmable to realize an arbitrary regular expression.

In this invention I also show how the content search processor of this invention can be used
20 to create general application acceleration in a compute device like a server, personal computer, workstation or the like. I show an example content search API which can be used as a general facility that may get offered by an operating system for those devices to applications running on them which can utilize the content search processor and significantly improve the performance of those applications compared to having them run on the general
25 purpose processor of these devices.

An example application of anti-spam is illustrated in this application which can be accelerated to become a high line rate application unlike current solutions which run on general purpose processor of the compute servers on which they run.

I also illustrate an example of using the content search processor coupled with a protocol
30 processor such that the content search processing can be done in line with the traffic and appropriate actions be taken dependent on the traffic content. The processor of this invention can thus be used to apply content specific search rules for traffic stream. For example, if a specific packet or stream of packets carry SMTP traffic, then the protocol processor can let the content search processor know that and provide the appropriate flow

information. Then, the content search processor retrieves the flow context for the current flow from the memory and retrieves the SMTP application rules context from the application memory associated with the search engines. The search engines get configured to process the content of the specific flow and its associated application context. Thus content search
5 can continue across multiple packets of the same flow even when packets in the flow arrive with significant time gap between them and multiple different application rule contexts can also be applied. By using such an architecture a significant performance benefit can result compared to architectures where the rule context can not be changed rapidly, and where the context may need to be brought in from global memory into each of the FSA engines.

10 I also describe a high performance hardware processor that sharply reduces the TCP/IP protocol stack overhead from host processor and enables a high line rate storage and data transport solution based on IP.

This patent also describes the novel high performance processor that sharply reduces the TCP/IP protocol stack overhead from the host processor and enables high line rate security
15 processing including firewall, encryption, decryption, intrusion detection and the like. This patent also describes a content inspection architecture that may be used for detecting spam, viruses, digital rights management information, instant message inspection, URL matching, application detection, malicious content, and other content and applying specific rules which may enable anti-spam, anti-virus and the like capabilities. The content inspection engine
20 may be used for detecting and enforcing digital rights management rules for the content. The content inspection engine may also be used for URL matching, string searches, content based load balancing, sensitive information search like credit card numbers or social security numbers or health information or the like. The content inspection engine results may be used to direct the operation of the run-time adaptable processor as well.

25 This patent also describes a novel processor architecture that is run-time adaptable to the needs of the data sent to or received from a network. The run-time adaptable features of this processor can be used to deploy services that operate on network data under control of user definable policies. The adaptable processor may also be used to dynamically offload compute intensive operations from the host processor, when not performing operations on
30 the network data or in conjunction with network data processing if enough adaptable hardware resources are available. The processor performs protocol processing like TCP/IP or SCTP or UDP or the like using the high performance protocol processor disclosed and then uses an adaptable processing hardware to provide other functions or services like socket layer security, Transport layer security, encryption/decryption, RDMA, RDMA
35 security, application layer security, content inspection, deep packet inspection, virus

scanning or detection, policy processing, content based switching, load balancing, content based load balancing, virtualization or higher application layer processing or a combination thereof. Higher layer processing may further involve application level protocol processing (for example, protocol processing for HTTP, HTTPS, XML, SGML, Secure XML, other XML derivatives, Telnet, FTP, IP Storage, NFS, CIFS, DAFS and the like) which may also be
5 accelerated by dynamically adapting or reconfiguring the processor of this patent. This can significantly reduce the processing overhead on the host processor of the target system, without adding major system cost of adding dedicated accelerator hardware.

The processing capabilities of a system deploying the runtime adaptable processor of this
10 patent can continue to expand and improve without the need for continually upgrading the system with host processor to achieve performance benefits. The hardware of the processor may comprise computational elements organized into compute clusters. Computational elements may provide logical and arithmetic operations beside other functions. A computational element may operate on 1-bit, 2-bit, 4-bit, 8-bit or n-bit data sizes as may be
15 chosen by the implementation. Thus multiple computational elements may together provide the desired size of operators. For example, if each computational element provides operations on a largest data size of 8-bits, then to operate on 32-bit operands, four computational elements may each operate on a byte of the 32-bit operand. The computational elements within the compute clusters can be programmatically interconnected
20 with each other using a dynamically changeable or adaptable interconnection network. The compute clusters may also be dynamically interconnected with each other programmatically forming an adaptable network. Thus arbitrary interconnections can be created between the array of computational elements within a compute cluster as well as outside of the compute clusters. The computational elements may each be dynamically changed to provide
25 necessary function(s) or operation(s) by programmatically selecting and connecting the necessary logic blocks through muxes or other logic blocks or other means. The computational elements may also be simple processors with ALU and other functional blocks. In this case to change hardware function of the computational element (CE), it is programmatically instructed to execute certain function(s) or operation(s). The operation(s)
30 selected for each of the computational element can be different. Thus the processor hardware of this patent can be dynamically i.e. during operation or at runtime, be changed or adapted to provide a different functionality. For explanation purposes let us take an example of a compute cluster with 8 computational elements each providing 8-bit operations. If we want to perform two 32-bit operations like a 32-bit addition followed by an n-bit shift
35 operation, then the computational elements may be grouped into two sets of four each. The first group would be programmed to provide addition operation, where each of them may

operate on 8-bits at a time. The appropriate carry and flags and other outputs would be available through the interconnections between the CEs which may be programmatically selected. The second group of CEs would be programmed to provide a shift operation on the incoming data. One such setup of the CEs may be called an Avatar or a virtual
5 configuration/setup. The CEs may then continue to provide these operations on the input operands for a period of time that this avatar is maintained. Then it is possible to dynamically change the avatar to a new avatar. For instance in the example used above, let us assume that after a certain time period, which may be as small as a clock period or multiple clock periods or other period, the processor needs to switch from providing
10 acceleration support from 32-bit Add, followed by Shift to something like two 16-bit Subtraction followed by two 16-bit logical AND. In such an instance the hardware is setup to form four groups of two CEs, each group operating on 16-bit operands. First four CEs in this case may now be dynamically switched or changed from providing addition function to subtraction function. Further, they may now be dynamically switched to operate in two
15 groups to provide 16-bit operation, instead of one group providing 32-bit operation in the previous avatar. Similarly, the second group of four CEs from the previous avatar may now be dynamically switched or changed to provide logical AND operations and may also be setup as two groups providing 16-bit operations. This forms a new avatar of the hardware which has been dynamically changed as per the need of the required functionality at the
20 time. Thus the runtime adaptable protocol processor of this patent can change its functions at a fine granularity along with the interconnections of these operators provided by the CEs to form runtime changeable or adaptable hardware platform. The operations supported may be lot more complex than those used in the examples discussed above. The examples were provided primarily to provide a better appreciation of the capabilities and were deliberately
25 kept simplistic. Though the examples suggested a unidirectional flow, this is not to be construed as the only mode of operation. The outputs from the operations in the examples above could be recycled to the first group of CEs which would allow a pipelined loop of the hardware. More complex scenarios are feasible with the dynamically adaptable nature of the CEs and the interconnection network, where different stages of CEs may be switched
30 over a period of time to provide different functionality as may be required by the algorithm or application or service being enabled in hardware. Pipelined stages of operations are thus possible with arbitrary loop backs as necessary. Hence the applications or services being accelerated or supported in hardware can increase over time, where the users may decide to accelerate applications of choice by mapping them appropriately to the runtime adaptable
35 protocol processor as and when necessary or feasible due to cost, performance, resources, application discovery, application development or any other reasons that may cause them to invent or develop new applications or services. The hardware system may thus be adapted

to use such applications without the need for incurring costs of buying new systems or accelerators and the like. The system capabilities can be increased over time as new services are developed and/or deployed that exploit the adaptable component of the processor of this invention. The new services or policies or a combination thereof may be
5 deployed to the appropriate systems over a network under user control.

Traditionally, TCP/IP networking stack is implemented inside the operating system kernel as a software stack. The software TCP/IP stack implementation consumes, as mentioned above, more than 50% of the processing cycles available in a 1 GHz processor when serving a 1Gbps network. The overhead comes from various aspects of the software
10 TCP/IP stack including checksum calculation, memory buffer copy, processor interrupts on packet arrival, session establishment, session tear down and other reliable transport services. The software stack overhead becomes prohibitive at higher lines rates. Similar issues occur in networks with lower line rates, like wireless networks, that use lower performance host processors. A hardware implementation can remove the overhead from
15 the host processor.

The software TCP/IP networking stack provided by the operating systems uses up a majority of the host processor cycles. TCP/IP is a reliable transport that can be run on unreliable data links. Hence, when a network packet is dropped or has errors, TCP does the retransmission of the packets. The errors in packets are detected using checksum that is
20 carried within the packet. The recipient of a TCP packet performs the checksum of the received packet and compares that to the received checksum. This is an expensive compute intensive operation performed on each packet involving each received byte in the packet. The packets between a source and destination may arrive out of order and the TCP layer performs ordering of the data stream before presenting it to the upper layers. IP
25 packets may also be fragmented based on the maximum transfer unit (MTU) of the link layer and hence the recipient is expected to de-fragment the packets. These functions result in temporarily storing the out of order packets, fragmented packets or unacknowledged packets in memory on the network card for example. When the line rates increase to above 1Gbps, the memory size overhead and memory speed bottleneck resulting from these add
30 significant cost to the network cards and also cause huge performance overhead. Another function that consumes a lot of processor resources is the copying of the data to/from the network card buffers, kernel buffers and the application buffers.

Microprocessors are increasingly achieving their high performance and speed using deep pipelining and super scalar architectures. Interrupting these processors on arrival of small
35 packets will cause severe performance degradation due to context switching overhead,

pipeline flushes and refilling of the pipelines. Hence interrupting the processors should be minimized to the most essential interrupts only. When the block storage traffic is transported over TCP/IP networks, these performance issues become critical, severely impacting the throughput and the latency of the storage traffic. Hence the processor intervention in the entire process of transporting storage traffic needs to be minimized for IP based storage solutions to have comparable performance and latency as other specialized network architectures like fibre channel, which are specified with a view to a hardware implementation. Emerging IP based storage standards like iSCSI, FCIP, iFCP, and others (like NFS, CIFS, DAFS, HTTP, XML, XML derivatives (such as Voice XML, EBXML, Microsoft SOAP and others), SGML, and HTML formats) encapsulate the storage and data traffic in TCP/IP segments. However, there usually isn't alignment relationship between the TCP segments and the protocol data units that are encapsulated by TCP packets. This becomes an issue when the packets arrive out of order, which is a very frequent event in today's networks. The storage and data blocks cannot be extracted from the out of order packets for use until the intermediate packets in the stream arrive which will cause the network adapters to store these packets in the memory, retrieve them and order them when the intermediate packets arrive. This can be expensive from the size of the memory storage required and also the performance that the memory subsystem is expected to support, particularly at line rates above 1Gbps. This overhead can be removed if each TCP segment can uniquely identify the protocol data unit and its sequence. This can allow the packets to be directly transferred to their end memory location in the host system. Host processor intervention should also be minimized in the transfer of large blocks of data that may be transferred to the storage subsystems or being shared with other processors in a clustering environment or other client server environment. The processor should be interrupted only on storage command boundaries to minimize the impact.

The IP processor set forth herein eliminates or sharply reduces the effect of various issues outlined above through innovative architectural features and the design. The described processor architecture provides features to terminate the TCP traffic carrying the storage and data payload thereby eliminating or sharply reducing the TCP/IP networking stack overhead on the host processor, resulting in packet streaming architecture that allows packets to pass through from input to output with minimal latency. To enable high line rate storage or data traffic being carried over IP requires maintaining the transmission control block information for various connections (sessions) that are traditionally maintained by host kernel or driver software. As used in this patent, the term "IP session" means a session for a session oriented protocol that runs on IP. Examples are TCP/IP, SCTP/IP, and the like. Accessing session information for each packet adds significant processing overhead. The

described architecture creates a high performance memory subsystem that significantly reduces this overhead. The architecture of the processor provides capabilities for intelligent flow control that minimizes interrupts to the host processor primarily at the command or data transfer completion boundary.

5 Today, no TCP/IP processor is offered with security.

The conventional network security model deployed today involves perimeter security in the form of perimeter firewall and intrusion detection systems. However, as increasing amount of business gets conducted on-line, there is a need to provide enterprise network access to "trusted insiders" - employees, partners, customers and contractors from outside. This
10 creates potential threats to the information assets inside an enterprise network. Recent research by leading firms and FBI found that over 70 per cent of the unauthorized access to information systems is committed by employees or trusted insiders and so are over 95 per cent of intrusions that result in substantial financial loss. In an environment where remote access servers, peer networks with partners, VPN and wireless access points blur the
15 boundary of the network, a perimeter security is not sufficient. In such an environment organizations need to adopt an integrated strategy that addresses network security at all tiers including at the perimeter, gateways, servers, switches, routers and clients instead of using point security products at the perimeter.

Traditional firewalls provide perimeter security at network layers by keeping offending IP
20 addresses out of the internal network. However, because many new attacks arrive as viruses or spam, exploiting known vulnerabilities of well-known software and higher level protocols, it is desirable to develop and deploy application layer firewalls. These should also be distributed across the network instead of being primarily at the perimeter.

Currently as the TCP/IP processing exists as the software stack in clients, servers and other
25 core and end systems, the security processing also is done in software particularly the capabilities like firewall, intrusion detection and prevention. As the line rates of these networks go to 1Gbps and 10Gbps, it is imperative that the TCP/IP protocol stack be implemented in hardware because a software stack consumes a large portion of the available host processor cycles. Similarly, if the security processing functions get deployed
30 on core or end systems instead of being deployed only at the perimeter, the processing power required to perform these operations may create a huge overhead on the host processor of these systems. Hence software based distributed security processing would increase the required processing capability of the system and increase the cost of deploying such a solution. A software based implementation would be detrimental to the performance

- of the servers and significantly increase the delay or latency of the server response to clients and may limit the number of clients that can be served. Further, if the host system software stack gets compromised during a network attack, it may not be possible to isolate the security functions, thereby compromising network security. Further, as the TCP/IP protocol
- 5 processing comes to be done in hardware, the software network layer firewalls may not have access to all state information needed to perform the security functions. Hence, the protocol processing hardware may be required to provide access to the protocol layer information that it processes and the host may have to redo some of the functions to meet the network firewall needs.
- 10 The hardware based TCP/IP and security rules processing processor of this patent solves the distributed core security processing bottleneck besides solving the performance bottleneck from the TCP/IP protocol stack. The hardware processor of this patent sharply reduces the TCP/IP protocol stack processing overhead from the host CPU and enables security processing features like firewall at various protocol layers such as link, network and
- 15 transport layers, thereby substantially improving the host CPU performance for intended applications. Further, this processor provides capabilities that can be used to perform deep packet inspection to perform higher layer security functions using the programmable processor and the classification/policy engines disclosed. This patent also describes a content inspection architecture that may be used for detecting spam, viruses, digital rights
- 20 management information, instant message inspection, URL matching, application detection, malicious content, and other content and applying specific rules which may enable anti-spam, anti-virus and the like security and content inspection and processing capabilities. The processor of this patent thus enables hardware TCP/IP and security processing at all layers of the OSI stack to implement capabilities like firewall at all layers including the
- 25 network layer and application layers.

- The processor architecture of this patent also provides integrated advanced security features. This processor allows for in-stream encryption and decryption of the network traffic on a packet by packet basis thereby allowing high line rates and at the same time offering confidentiality of the data traffic. Similarly, when the storage traffic is carried on a network
- 30 from the server to the storage arrays in a SAN or other storage system, it is exposed to various security vulnerabilities that a direct attached storage system does not have to deal with. This processor allows for in stream encryption and decryption of the storage traffic thereby allowing high line rates and at the same time offering confidentiality of the storage data traffic.

Classification of network traffic is another task that consumes up to half of the processing cycles available on packet processors leaving few cycles for deep packet inspection and processing. IP based storage traffic by the nature of the protocol requires high speed low latency deep packet processing. The described IP processor significantly reduces the classification overhead by providing a programmable classification engine. The programmable classification engine of this patent allows deployment of advanced security policies that can be enforced on a per packet, per transaction, and per flow basis. This will result in significant improvement in deploying distributed enterprise security solutions in a high performance and cost effective manner to address the emerging security threats from within the organizations.

To enable the creation of distributed security solutions, it is critical to address the need of Information Technology managers to cost effectively manage the entire network. Addition of distributed security, without means for ease of managing it can significantly increase the management cost of the network. The disclosure of this patent also provides a security rules/policy management capability that can be used by IT personnel to distribute the security rules from a centralized location to various internal network systems that use the processor of this patent. The processor comprises hardware and software capabilities that can interact with centralized rules management system(s). Thus the distribution of the security rules and collection of information of compliance or violation of the rules or other related information like offending systems, users and the like can be processed from one or more centralized locations by IT managers. Thus multiple distributed security deployments can be individually controlled from centralized location(s).

This patent also provides means to create a secure operating environment for the protocol stack processing that, even if the host system gets compromised either through a virus or malicious attack, allows the network security and integrity to be maintained. This patent significantly adds to the trusted computing environment needs of the next generation computing systems.

Tremendous growth in the storage capacity and storage networks have created storage area management as a major cost item for IT departments. Policy based storage management is required to contain management costs. The described programmable classification engine allows deployment of storage policies that can be enforced on packet, transaction, flow and command boundaries. This will have significant improvement in storage area management costs.

The programmable IP processor architecture also offers enough headroom to allow customer specific applications to be deployed. These applications may belong to multiple categories e.g. network management, storage firewall or other security capabilities, bandwidth management, quality of service, virtualization, performance monitoring, zoning,

5 LUN masking and the like.

The adaptable processor hardware may be used to accelerate many of the applications or services listed above based on the available reprogrammable resources, deployed applications, services, policies or a combination thereof.

BRIEF DESCRIPTION OF THE DRAWINGS

10 Fig. 1 illustrates a layered SCSI architecture and interaction between respective layers located between initiator and target systems.

Fig. 2 illustrates the layered SCSI architecture with iSCSI and TCP/IP based transport between initiator and target systems.

Fig. 3 illustrates an OSI stack comparison of software based TCP/IP stack with hardware -
15 oriented protocols like Fibre channel.

Fig. 4 illustrates an OSI stack with a hardware based TCP/IP implementation for providing performance parity with the other non-IP hardware oriented protocols.

Fig. 5 illustrates a host software stack illustrating operating system layers implementing networking and storage stacks.

20 Fig. 6 illustrates software TCP stack data transfers.

Fig. 7 illustrates remote direct memory access data transfers using TCP/IP offload from the host processor as described in this patent.

Fig. 8 illustrates host software SCSI storage stack layers for transporting block storage data over IP networks.

25 Fig. 9 illustrates certain iSCSI storage network layer stack details of an embodiment of the invention.

Fig. 10 illustrates TCP/IP network stack functional details of an embodiment of the invention.

Fig. 11 illustrates an iSCSI storage data flow through various elements of an embodiment of the invention.

Fig. 12 illustrates iSCSI storage data structures useful in the invention.

Fig. 13 illustrates a TCP/IP Transmission Control Block data structure for a session
5 database entry useful in an embodiment of the invention.

Fig. 14 illustrates an iSCSI session database structure useful in an embodiment of the invention.

Fig. 15 illustrates iSCSI session memory structure useful in an embodiment of the invention.

Fig. 16 illustrates a high-level architectural block diagram of an IP network application
10 processor useful in an embodiment of the invention.

Fig. 17 illustrates a detailed view of the architectural block diagram of the IP network application processor of Fig. 16.

Fig. 18 illustrates an input queue and controller for one embodiment of the IP processor.

Fig. 19 illustrates a packet scheduler, sequencer and load balancer useful in one
15 embodiment of the IP processor.

Fig. 20 illustrates a packet classification engine, including a policy engine block of one embodiment of the IP storage processor.

Fig. 21 broadly illustrates an embodiment of the SAN packet processor block of one embodiment of an IP processor at a high-level.

Fig. 22 illustrates an embodiment of the SAN packet processor block of the described IP
20 processor in further detail.

Fig. 23 illustrates an embodiment of the programmable TCP/IP processor engine which can be used as part of the described SAN packet processor.

Fig. 24 illustrates an embodiment of the programmable IP Storage processor engine which
25 can be used as part of the described SAN packet processor.

Fig. 25 illustrates an embodiment of an output queue block of the programmable IP processor of Fig. 17.

Fig. 26 illustrates an embodiment of the storage flow controller and RDMA controller.

Fig. 27 illustrates an embodiment of the host interface controller block of the IP processor useful in an embodiment of the invention.

Fig. 28 illustrates an embodiment of the security engine.

- 5 Fig. 29 illustrates an embodiment of a memory and controller useful in the described processor.

Fig. 30 illustrates a data structure useable in an embodiment of the described classification engine.

Fig. 31 illustrates a storage read flow between initiator and target.

- 10 Fig. 32 illustrates a read data packet flow through pipeline stages of the described processor.

Fig. 33 illustrates a storage write operation flow between initiator and target.

Fig. 34 illustrates a write data packet flow through pipeline stages of the described processor.

- 15 Fig. 35 illustrates a storage read flow between initiator and target using the remote DMA (RDMA) capability between initiator and target.

Fig. 36 illustrates a read data packet flow between initiator and target using RDMA through pipeline stages of the described processor.

Fig. 37 illustrates a storage write flow between initiator and target using RDMA capability.

- 20 Fig. 38 illustrates a write data packet flow using RDMA through pipeline stages of the described processor.

Fig. 39 illustrates an initiator command flow in more detail through pipeline stages of the described processor.

- 25 Fig. 40 illustrates a read packet data flow through pipeline stages of the described processor in more detail.

Fig. 41 illustrates a write data flow through pipeline stages of the described processor in more detail.

Fig. 42 illustrates a read data packet flow when the packet is in cipher text or is otherwise a secure packet through pipeline stages of the described processor.

Fig. 43 illustrates a write data packet flow when the packet is in cipher text or is otherwise a secure packet through pipeline stages of the described processor of one embodiment of the
5 invention.

Fig. 44 illustrates a RDMA buffer advertisement flow through pipeline stages of the described processor.

Fig. 45 illustrates a RDMA write flow through pipeline stages of the described processor in more detail.

10 Fig. 46 illustrates a RDMA Read data flow through pipeline stages of the described processor in more detail.

Fig. 47 illustrates steps of a session creation flow through pipeline stages of the described processor.

Fig. 48 illustrates steps of a session tear down flow through pipeline stages of the described
15 processor.

Fig. 49 illustrates a session creation and session teardown steps from a target perspective through pipeline stages of the described processor.

Fig. 50 illustrates an R2T command flow in a target subsystem through pipeline stages of the described processor.

20 Fig. 51 illustrates a write data flow in a target subsystem through pipeline stages of the described processor.

Fig. 52 illustrates a target read data flow through the pipeline stages of the described processor.

Fig. 53 illustrates a typical enterprise network with perimeter security.

25 Fig. 54 illustrates an enterprise network with distributed security using various elements of this patent.

Fig. 55 illustrates an enterprise network with distributed security including security for a storage area network using various elements of this patent.

- Fig. 56 illustrates a Central Manager/Policy Server & Monitoring Station.
- Fig. 57 illustrates Central Manager flow of the disclosed security feature.
- Fig. 58 illustrates rule distribution flow for the Central Manager.
- Fig. 59 illustrates Control Plane Processor/Policy Driver Flow for the processor of this patent.
- 5 Fig. 60 illustrates a sample of packet filtering rules that may be deployed in distributed security systems.
- Fig. 61 illustrates a TCP/IP processor version of the IP processor of Figs. 16 and 17.
- Fig. 62 illustrates an adaptable TCP/IP processor.
- Fig. 63 illustrates an adaptable TCP/IP processor useable as an alternate to that of Fig. 62.
- 10 Fig. 64 illustrates a runtime adaptable processor.
- Fig. 65 illustrates a compute cluster.
- Fig. 66 illustrates a security solution for providing security in a network
- Fig. 67 illustrates a security solution compiler flow
- Fig. 68 illustrates a flow-through secure network card architecture
- 15 Fig. 69 illustrates a network line card with look-aside security architecture
- Fig. 70 illustrates a security and content search accelerator adapter architecture
- Fig. 71 illustrates a security processor architecture
- Fig. 72 illustrates an alternate security processor architecture
- Fig. 73 illustrates a third security processor architecture
- 20 Fig. 74 illustrates a content search and rule processing engine architecture
- Fig. 75 illustrates an example of Regular Expression Rules (prior art)
- Fig. 76a illustrates Thompson's NFA (prior art)
- Fig. 76b illustrates Berry-Sethi NFA (prior art)

- Fig. 76c illustrates DFA (prior art)
- Fig. 77a illustrates a left-biased NFA and state transition table (prior art)
- Fig. 77b illustrates a right-biased NFA and state transition table (prior art)
- Fig. 78a illustrates state transition controls
- 5 Fig. 78b illustrates configurable next state tables per state
- Fig. 79a illustrates state transition logic (STL) for a state
- Fig. 79b illustrates a state logic block
- Fig. 80 illustrates a NFA based search engine
- Fig. 81 illustrates application state memory example configuration
- 10 Fig. 82 illustrates a DFA based search engine
- Fig. 83 illustrates example DFA operations
- Fig. 84 illustrates a runtime adaptable search processor
- Fig. 85 illustrates a runtime adaptable search processor array
- Fig. 86 illustrates an example search cluster
- 15 Fig. 87 illustrates Intra search cluster rule groups and switching example
- Fig. 88 illustrates an example rules configuration and context switch
- Fig. 89 illustrates a computing device with content search accelerator
- Fig. 90 illustrates an example anti-spam performance bottleneck and solution
- Fig. 91 illustrates anti-spam with anti-virus performance bottleneck
- 20 Fig. 92 illustrates application content search performance bottleneck and solution
- Fig. 93 illustrates an example content search API usage model
- Fig. 94 illustrates an example content search API with example functions
- Fig. 95 illustrates an example application flow (static setup) using search processor

Fig. 96 illustrates FSA compiler flow

Fig. 97 illustrates search compiler flow (full + incremental rule distribution)

Fig. 98 illustrates FSA Synthesis & Compiler Flow for FPGA

DESCRIPTION

- 5 I provide a new high performance content search processor that can relieve the performance bottleneck of content search from the host processor. The search processor of this patent is used to perform a large number of regular expression searches in parallel using NFA and DFA based search engines. The search processor can perform high speed content search at line rates from 1Gbps to 10Gbps and higher, when the best of class server
- 10 microprocessor can only perform the same tasks at well below 100Mbps. The content search processor can be used not only to perform layer 2 through layer 4 searches that may be used for classification and security applications, it can also be used to perform deep packet inspection and layer 4 through layer 7 content analysis and security applications.

- I provide a new high performance and low latency way of implementing a TCP/IP stack in
- 15 hardware to relieve the host processor of the severe performance impact of a software TCP/IP stack. This hardware TCP/IP stack is then interfaced with additional processing elements to enable high performance and low latency IP based storage applications.

- This system also enables a new way of implementing security capabilities like firewall inside enterprise networks in a distributed manner using a hardware TCP/IP implementation with
- 20 appropriate security capabilities in hardware having processing elements to enable high performance and low latency IP based network security applications. The hardware processors may be used inside network interface cards of servers, workstations, client PCs, notebook computers, handheld devices, switches, routers and other networked devices. The servers may be web servers, remote access servers, file servers, departmental servers,
- 25 storage servers, network attached storage servers, database servers, blade servers, clustering servers, application servers, content /media servers, grid computers/servers, and the like. The hardware processors may also be used inside an I/O chipset of one of the end systems or network core systems like a switch or router or appliance or the like.

- This system enables distributed security capabilities like firewall, intrusion detection, virus
- 30 scan, virtual private network, confidentiality services and the like in internal systems of an enterprise network. The distributed security capabilities may be implemented using the hardware processor of this patent in each system, or some of its critical systems and others

may deploy those services in software. Hence, overall network will include distributed security as hardware implementation or software implementation or a combination thereof in different systems depending on the performance, cost and security needs as determined by IT managers. The distributed security systems will be managed from one or more
5 centralized systems used by IT managers for managing the network using the principles described. This will enable an efficient and consistent deployment of security in the network using various elements of this patent.

This can be implemented in a variety of forms to provide benefits of TCP/IP termination, high performance and low latency IP storage capabilities, remote DMA (RDMA) capabilities,
10 security capabilities, programmable classification, policy processing features, runtime adaptable processing, high speed content search and the like. Following are some of the embodiments that can implement this:

Server

The described architecture may be embodied in a high performance server environment
15 providing hardware based TCP/IP functions or hardware TCP/IP and security or search functions or a combination of the foregoing that relieve the host server processor or processors of TCP/IP and/or security and/or search software and their performance overhead. Any processor of this invention may be a companion processor to a server
chipset, providing the high performance networking interface with hardware TCP/IP and/or
20 security and/or content search. Servers can be in various form factors like blade servers, appliance servers, file servers, thin servers, clustered servers, database server, content search server, game server, grid computing server, VOIP server, wireless gateway server, security server, network attached storage server or traditional servers. The current
embodiment would allow creation of a high performance network interface on the server
25 motherboard.

Further, the described runtime adaptable protocol processor or security processor or search processor or a combination of the foregoing architectures may also be used to provide additional capabilities or services beside protocol processing like socket layer security, Transport layer security, encryption/decryption, RDMA, RDMA security, application layer
30 security, virtualization or higher application layer processing which may further involve application level protocol processing (for example, protocol processing for HTTP, HTTPS, XML, SGML, Secure XML, other XML derivatives, Telnet, FTP, IP Storage, NFS, CIFS, DAFS, and the like). One embodiment could include TCP/IP protocol processing using the dedicated protocol processor and XML acceleration mapped to a runtime adaptable

processor such as that disclosed in this patent. The protocol processor may or may not provide RDMA capabilities dependent upon the system needs and the supported line rates. Security processing capabilities of this invention may also be optionally incorporated in this embodiment. The same architecture could also be used to provide security acceleration support to XML data processing on the runtime adaptable processor of this patent.

Companion Processor to a server Chipset

The server environment may also leverage the high performance IP storage processing capability of the described processor, besides high performance TCP/IP and/or RDMA capabilities. In such an embodiment the processor may be a companion processor to a server chipset providing high performance network storage I/O capability besides the TCP/IP offloading from the server processor. This embodiment would allow creation of high performance IP based network storage I/O on the motherboard. In other words it would enable IP SAN on the motherboard.

Similar to the Server embodiment described above, this embodiment may also leverage the runtime adaptable processor of this patent to provide adaptable hardware acceleration along with protocol processing support in a server chipset. The runtime adaptable processor can be configured to provide storage services like virtualization, security services, multi-pathing, protocol translation and the like. Protocol translation may included for example translation to/from FibreChannel protocol to IP Storage protocol or vice versa, Serial ATA protocol to IP Storage or FibreChannel protocol or vice-versa, Serial Attached SCSI protocol to IP Storage or FibreChannel protocol or vice-versa, and the like.

Further, the runtime adaptable search processor of this patent may provide content search acceleration capability to a server chipset similar to the server embodiment described above. The content search processor may be used for a vast variety of applications that use regular expressions for the search rules. The applications cover a spectrum of areas like bioinformatics, application and network layer security, proteomics, genetics, web search, protocol analysis, XML acceleration, data warehousing and the like.

Storage System Chipsets

The processor may also be used as a companion of a chipset in a storage system, which may be a storage array (or some other appropriate storage system or subsystem) controller, which performs the storage data server functionality in a storage networking environment. The processor would provide IP network storage capability to the storage array controller to network in an IP based SAN. The configuration may be similar to that in a server

environment, with additional capabilities in the system to access the storage arrays and provide other storage-centric functionality.

This embodiment may also leverage the runtime adaptable processor of this patent to provide adaptable hardware acceleration along with protocol processing support in a storage system chipset. The runtime adaptable processor can be configured to provide storage services like virtualization, security services, multi-pathing, protocol translation and the like. Protocol translation may included for example translation to/from FibreChannel protocol to IP Storage protocol or vice versa, Serial ATA protocol to IP Storage or FibreChannel protocol or vice-versa, Serial Attached SCSI protocol to IP Storage or FibreChannel protocol or vice-versa, and the like. The runtime adaptable processor may also be used to provide acceleration for storage system metadata processing to improve the system performance.

This embodiment may also leverage the runtime adaptable search processor of this patent to provide high speed search of storage traffic for various applications like confidential information policy enforcement, virus/malware detection, high speed indexing for ease of searching, virtualization, content based switching, security services and the like.

Server/Storage Host Adapter Card

The IP processor may also be embedded in a server host adapter card providing high speed TCP/IP networking. The same adapter card may also be able to offer high speed network security capability for IP networks. Similarly, the adapter card may also be able to offer high speed network storage capability for IP based storage networks. The adapter card may be used in traditional servers and may also be used as blades in a blade server configuration. The processor may also be used in adapters in a storage array (or other storage system or subsystem) front end providing IP based storage networking capabilities. The adapter card may also leverage the runtime adaptable processors of this patent in a way similar to that described above.

Processor Chipset Component

The TCP/IP processor may be embodied inside a processor chipset, providing the TCP/IP offloading capability. Such a configuration may be used in the high end servers, workstations or high performance personal computers that interface with high speed networks. Such an embodiment could also include IP storage or RDMA capabilities or combination of this invention to provide IP based storage networking and/or TCP/IP with RDMA capability embedded in the chipset. The usage of multiple capabilities of the described architecture can be made independent of using other capabilities in this or other

embodiments, as a trade-off of feature requirements, development timeline and cost, silicon die cost, and the like. The processor chipset may also incorporate the runtime adaptable processors of this patent to offer a variable set of functions on demand by configuring the processor for the desired application.

5 **Storage or SAN System or Subsystem Switching Line Cards**

The IP processor may also be used to create high performance, low latency IP SAN switching system (or other storage system or subsystem) line cards. The processor may be used as the main processor terminating and originating IP-based storage traffic to/from the line card. This processor would work with the switching system fabric controller, which may
10 act like a host, to transport the terminated storage traffic, based on their IP destination, to the appropriate switch line card as determined by the forwarding information base present in the switch system. Such a switching system may support purely IP based networking or may support multi-protocol support, allow interfacing with IP based SAN along with other data center SAN fabrics like Fibre channel. A very similar configuration could exist inside a
15 gateway controller system, that terminates IP storage traffic from LAN or WAN and originates new sessions to carry the storage traffic into a SAN, which may be IP based SAN or more likely a SAN built from other fabrics inside a data center like Fibre channel. The processor could also be embodied in a SAN gateway controller. These systems would use security capabilities of this processor to create a distributed security network within
20 enterprise storage area networks as well.

The runtime adaptable processor of this patent can be very effective in providing hardware acceleration capabilities individually or in combination as described above like protocol translation, virtualization, security, bandwidth management, rate limiting, grooming, policy based management and the like

25 The runtime adaptable search processor of this patent can be used to provide deep content inspection capability in this embodiment which can be used to created intelligent networking capabilities like content based switching, application oriented networking and the like.

Network Switches, routers, wireless access points

The processor may also be embedded in a network interface line card providing high speed
30 TCP/IP networking for switches, routers, gateways, wireless access points and the like. The same adapter card may also be able to offer high speed network security capability for IP networks. This processor would provide the security capabilities that can then be used in a distributed security network.

The runtime adaptable processor of this patent may also be used in such embodiments offering services and capabilities described above as well as others like Wired Equivalent Privacy security capabilities, RADIUS and like security features as needed by the environment. The runtime adaptable processor may also be used to provide dynamically changeable protocol processing capability besides TCP/IP processing to support wireless protocols like Bluetooth, HomeRF, wireless Ethernet LAN protocols at various line rates, 3GPP, GPRS, GSM, or other wireless LAN or RF or cellular technology protocols, or any combinations thereof.

The runtime adaptable search processor of this patent can be used to provide deep content inspection capability in this embodiment which can be used to create intelligent networking capabilities like content based switching, application oriented networking and the like.

Storage Appliance

Storage networks management costs are increasing rapidly. The ability to manage the significant growth in the networks and the storage capacity would require creating special appliances which would be providing the storage area management functionality. The described management appliances for high performance IP based SAN, would implement my high performance IP processor, to be able to perform its functions on the storage traffic transported inside TCP/IP packets. These systems would require a high performance processor to do deep packet inspection and extract the storage payload in the IP traffic to provide policy based management and enforcement functions. The security, programmable classification and policy engines along with the high speed TCP/IP and IP storage engines described would enable these appliances and other embodiments described in this patent to perform deep packet inspection and classification and apply the policies that are necessary on a packet by packet basis at high line rates at low latency. Further these capabilities can enable creating storage management appliances that can perform their functions like virtualization, policy based management, security enforcement, access control, intrusion detection, bandwidth management, traffic shaping, quality of service, anti-spam, virus detection, encryption, decryption, LUN masking, zoning, link aggregation and the like in-band to the storage area network traffic. Similar policy based management, and security operations or functionality may also be supported inside the other embodiments described in this patent. The runtime adaptable processor of this patent can be used to dynamically support or accelerate one or more of the applications/services. The services/applications supported may be selected by the policies in existence under the influence or control of the user or the administrator. The runtime adaptable search processor of this invention can be used to perform high speed network line rate searches in this embodiment which can be

used for several deep packet inspection, content search and security applications discussed above.

Clustered Environments

Server systems are used in a clustered environment to increase the system performance and scalability for applications like clustered data bases and the like. The applications running on high performance cluster servers require ability to share data at high speeds for inter-process communication. Transporting this inter-process communication traffic on a traditional software TCP/IP network between cluster processors suffers from severe performance overhead. Hence, specialized fabrics like Fibre channel have been used in such configurations. However, a TCP/IP based fabric which can allow direct memory access between the communicating processes' memory, can be used by applications that operate on any TCP/IP network without being changed to specialized fabrics like fibre channel. The described IP processor with its high performance TCP/IP processing capability and the RDMA features, can be embodied in a cluster server environment to provide the benefits of high performance and low latency direct memory to memory data transfers. This embodiment may also be used to create global clustering and can also be used to enable data transfers in grid computers and grid networks. The processor of this patent may also be used to accelerate local cluster data transfers using light weight protocols other than TCP/IP to avoid the latency and protocol processing overhead. The runtime adaptable processor architecture can be leveraged to support such a light weight protocol. Thus the same processor architecture may be used for local as well as global clustering and enable data transfers in grid computers and grid networks. The programmable processor of this patent may also be used for similar purposes without burdening the runtime adaptable processor. The processor architecture of this patent can thus be used to enable utility computing. The runtime adaptable processors of this patent may also be used to provide the capabilities described in other embodiments above to the clustered environment as well.

XML Accelerator

The runtime adaptable TCP/IP processor of this patent can also be used as a component inside a system or adapter card or as part of a chipset providing TCP/IP protocol termination or XML acceleration or a combination thereof. As web services usage increases, more and more web documents may start using XML or XML derivatives. The burden of processing XML on each web page access can be very significant on the host processors, requiring additional hardware support. The runtime adaptable processor of this patent can be used in such an environment to provide acceleration to XML processing, whereas transport protocol

processing is handled by the dedicated protocol processor of this patent. XML documents may also need security support, in which case the processor can be dynamically configured to provide security acceleration for secure XML documents. Similarly, the runtime adaptable search processor of this invention can also be used in such an embodiment, and provide

5 high speed XML processing capability to meet 1Gbps, 10Gbps and higher line rates. The search processor of this invention allows multiple thousand regular expression rules to be simultaneously evaluated for this application and can thus be used in networking systems also to provide intelligent networking and application oriented networking capabilities.

Voice Over IP (VOIP) Appliances

10 The processor of this patent can also be embedded inside voice over IP appliances like VOIP phones, servers, gateways, handheld devices, and the like. The protocol processor can be used to provide IP protocol processing, as well as the transport layer protocol processing as needed in the VOIP environment. Further, the runtime adaptable processor may be dynamically adapted to provide signal processing and DSP hardware acceleration

15 capabilities that may be required for VOIP appliance and the applications running on the appliance. The runtime adaptable search processor of this application can be used for speech recognition applications beside the VOIP applications similar to the discussion above.

Handheld Devices

20 The processor of this patent may also be used to provide protocol processing hardware capability to processors or chipsets of handheld devices, phones, personal digital assistants and the like. The protocol processor along with the runtime adaptable processor may provide many of the capabilities described above for many of the embodiments. The processor of this patent may be used to create a secure protocol processing stack inside

25 these devices as well as provide other services using hardware acceleration. The runtime adaptable processor may be used to enable the handheld devices to network in a wired or wireless manner. The device can then be dynamically adapted to work with a multitude of protocols like Bluetooth, Wireless Ethernet LAN, RF, GPRS, GSM, CDMA, CMDA variants or other 3G cellular technology or other wireless or cellular or RF technologies by using the

30 protocol processor and the runtime adaptable processor of this patent. The runtime adaptable search processor of this patent may also be embodied in handheld devices to accelerate the search performance of the already low performance processors. The search processor of this application may be incorporated in a chipset or a processor of the handheld device to accelerate the search performance of the applications running on these platforms.

The search processor may also be used as a multi-purpose security processor offering cryptographic security, network layer security functions through application layer security that may be embodied in a chipset or the host processor or as a companion processor in the handheld devices to protect against the emerging security threats for a variety of handheld devices.

Additional Embodiments

The processor architecture can be partially implemented in software and partially in hardware. The performance needs and cost implications can drive trade-offs for hardware and software partitioning of the overall system architecture of this invention. It is also possible to implement this architecture as a combination of chip sets along with the hardware and software partitioning or independent of the partitioning. For example the security processor and the classification engines could be on separate chips and provide similar functions. This can result in lower silicon cost of the IP processor including the development and manufacturing cost, but it may in some instances increase the part count in the system and may increase the footprint and the total solution cost. Security and classification engines could be separate chips as well. As used herein, a chip set may mean a multiple-chip chip set, or a chip set that includes only a single chip, depending on the application.

The storage flow controller and the queues could be maintained in software on the host or may become part of another chip in the chipset. Hence, multiple ways of partitioning this architecture are feasible to accomplish the high performance IP based storage and TCP/IP offload applications that will be required with the coming high performance processors in the future. The storage engine description has been given with respect to iSCSI, however, with TCP/IP and storage engine programmability, classifier programmability and the storage flow controller along with the control processor, other IP storage protocols like iFCP, FCIP and others can be implemented with the appropriate firmware. iSCSI operations may also represent IP Storage operations. The high performance IP processor core may be coupled with multiple input output ports of lower line rates, matching the total throughput to create multi-port IP processor embodiment as well.

It is feasible to use this architecture for high performance TCP/IP offloading from the main processor without using the storage engines. This can result in a silicon and system solution for next generation high performance networks for the data and telecom applications. The TCP/IP engine can be augmented with application specific packet accelerators and leverage the core architecture to derive new flavors of this processor. It is possible to change the

storage engine with another application specific accelerator like a firewall engine or a route look-up engine or a telecom/network acceleration engine, along with the other capabilities of this invention and target this processor architecture for telecom/networking and other applications.

- 5 The runtime adaptable search processor of this invention may also be embodied inside a variety of application specific systems targeted towards security or content search application or the like. For example, the search processor may be embodied inside an email security appliance that provides capabilities like anti-spam, anti-virus, anti-spyware, anti-worms, other malware prevention, as well as confidential information security and other
- 10 regulatory compliance for regulations like Sarbanes-Oxley, Gramm-Leach-Bliley, HIPAA and the like. The search processor may also be embodied in a system that may be used for extrusion detection and prevention systems to protect against confidential information leaks and other regulatory compliance like those listed above at an enterprise edge.

Detailed Description

- 15 Ability to perform content search has become a critical capability in the networked world. As the network line rates go up to 1Gbps, 10Gbps and higher, it is important to be able to perform deep packet inspection for many applications at line rate. Several security issues, like viruses, worms, confidential information leaks and the like, can be detected and prevented from causing damage if the network traffic can be inspected at high line rates. In
- 20 general, content search rules can be represented using regular expressions. Regular expression rules can be represented and computed using FSAs. NFAs and DFAs are the two types of FSAs that are used for evaluation of regular expressions. For high line rate applications a composite DFA can be used, where each character of the input stream can be processed per cycle of memory access. However, this does have a limit on how fast the
- 25 search can be performed dictated by the memory access speed. Another limiter of such approach is the amount of memory required to search even a modest number of regular expression rules. As discussed above, NFAs also have their limitations to achieve high performance on general purpose processors. In general, today's best of class microprocessors can only achieve less than 100Mbps performance using NFAs or DFAs for
- 30 a small number of regular expressions. Hence, there is a clear need to create targeted content search processor hardware to raise the performance of the search to the line rates of 1Gbps and 10Gbps. This invention shows such a high performance content search processor that can be targeted for high line rates.

Storage costs and demand have been increasing at a rapid pace over the last several years. This is expected to grow at the same rate in the foreseeable future. With the advent of e-business, availability of the data at any time and anywhere irrespective of the server or system downtime is critical. This is driving a strong need to move the server attached
5 storage onto a network to provide storage consolidation, availability of data and ease of management of the data. The storage area networks (SANs) are today predominantly based on Fibre Channel technology, that provide various benefits like low latency and high performance with its hardware oriented stacks compared to TCP/IP technology.

Some system transport block storage traffic on IP designed to transport data streams. The
10 data streams are transported using Transmission Control Protocol (TCP) that is layered to run on top of IP. TCP/IP is a reliable connection oriented protocol implemented in software within the operating systems. A TCP/IP software stack is slow to handle the high line rates that will be deployed in the future. New hardware solutions will accelerate the TCP/IP stack to carry storage and network traffic and be competitive to FC based solutions.

The prevalent storage protocol in high performance servers, workstations and storage
15 controllers and arrays is SCSI protocol which has been around for 20 years. SCSI architecture is built as layered protocol architecture. Fig. 1 illustrates the various SCSI architecture layers within an initiator, block 101, and target subsystems, block 102. As used in patent, the terms "initiator" and "target" mean a data processing apparatus, or a
20 subsystem or system including them. The terms "initiator" and "target" can also mean a client or a server or a peer. Likewise, the term "peer" can mean a peer data processing apparatus, or a subsystem or system thereof. A "remote peer" can be a peer located across the world or across the room.

The initiator and target subsystems in Fig. 1 interact with each other using the SCSI
25 application protocol layer, block 103, which is used to provide a client-server request and response transactions. It also provides device service request and response between the initiator and the target mass storage device which may take many forms like a disk arrays, tape drives, and the like. Traditionally, the target and initiator are interconnected using the SCSI bus architecture carrying the SCSI protocol, block 104. The SCSI protocol layer is the
30 transport layer that allows the client and the server to interact with each other using the SCSI application protocol. The transport layer must present the same semantics to the upper layer so that the upper layer protocols and application can stay transport protocol independent.

Fig. 2 illustrates the SCSI application layer on top of IP based transport layers. An IETF standards track protocol, iSCSI (SCSI over IP) is an attempt to provide IP based storage transport protocol. There are other similar attempts including FCIP (FC encapsulated in IP), iFCP(FC over IP) and others. Many of these protocols layer on top of TCP/IP as the
5 transport mechanism, in a manner similar to that illustrated in Fig. 2. As illustrated in Fig. 2, the iSCSI protocol services layer, block 204, provides the layered interface to the SCSI application layer, block 203. iSCSI carries SCSI commands and data as iSCSI protocol data units (PDUs) as defined by the standard. These protocol data units then can be transported over the network using TCP/IP, block 205, or the like. The standard does not specify the
10 means of implementing the underlying transport that carries iSCSI PDUs. Fig. 2 illustrates iSCSI layered on TCP/IP which provides the transport for the iSCSI PDUs.

The IP based storage protocol like iSCSI can be layered in software on top of a software based TCP/IP stack. However, such an implementation would suffer serious performance penalties arising from software TCP/IP and the storage protocol layered on top of that. Such
15 an implementation would severely impact the performance of the host processor and may make the processor unusable for any other tasks at line rates above 1Gbps. Hence, we would implement the TCP/IP stack in hardware, relieving the host processor, on which the storage protocol can be built. The storage protocol, like iSCSI, can be built in software running on the host processor or may, as described in this patent, be accelerated using
20 hardware implementation. A software iSCSI stack will present many interrupts to the host processor to extract PDUs from received TCP segments to be able to act on them. Such an implementation will suffer severe performance penalties for reasons similar to those for which a software based TCP stack would. The described processor provides a high performance and low latency architecture to transport Storage protocol on a TCP/IP based
25 network that eliminates or greatly reduces the performance penalty on the host processor, and the resulting latency impact.

Fig. 3 illustrates a comparison of the TCP/IP stack to Fibre channel as referenced to the OSI networking stack. The TCP/IP stack, block 303, as discussed earlier in the Summary of the Invention section of this patent, has performance problems resulting from the software
30 implementation on the hosts. Compared to that, specialized networking protocols like Fibre channel, block 304, and others are designed to be implemented in hardware. The hardware implementation allows the networking solutions to be higher performance than the IP based solution. However, the ubiquitous nature of IP and the familiarity of IP from the IT users' and developers' perspective makes IP more suitable for wide spread deployment. This can be
35 accomplished if the performance penalties resulting from TCP/IP are reduced to be

equivalent to those of the other competing specialized protocols. Fig. 4 illustrates a protocol level layering in hardware and software that is used for TCP/IP, block 403, to become competitive to the other illustrated specialized protocols.

Fig. 5 illustrates a host operating system stack using a hardware based TCP/IP and storage protocol implementation of this patent. The protocol is implemented such that it can be introduced into the host operating system stack, block 513, such that the operating system layers above it are unchanged. This allows the SCSI application protocols to operate without any change. The driver layer, block 515, and the stack underneath for IP based storage interface, block 501, will represent a similar interface as a non-networked SCSI interface, blocks 506 and 503 or Fibre Channel interface, block 502.

Fig. 6 illustrates the data transfers involved in a software TCP/IP stack. Such an implementation of the TCP/IP stack carries huge performance penalties from memory copy of the data transfers. The figure illustrates data transfer between client and server networking stacks. User level application buffers, block 601, that need to be transported from the client to the server or vice versa, go through the various levels of data transfers shown. The user application buffers on the source get copied into the OS kernel space buffers, block 602. This data then gets copied to the network driver buffers, block 603, from where it gets DMA-transferred to the network interface card (NIC) or the host bus adapter (HBA) buffers, block 604. The buffer copy operations involve the host processor and use up valuable processor cycles. Further, the data being transferred goes through checksum calculations on the host using up additional computing cycles from the host. The data movement into and out of the system memory on the host multiple times creates a memory bandwidth bottleneck as well. The data transferred to the NIC/HBA is then sent on to the network, block 609, and reaches the destination system. At the destination system the data packet traverses through the software networking stack in the opposite direction as the host though following similar buffer copies and checksum operations. Such implementation of TCP/IP stack is very inefficient for block storage data transfers and for clustering applications where a large amount of data may be transferred between the source and the destination.

Fig. 7 illustrates the networking stack in an initiator and in a target with features that allow remote direct memory access (RDMA) features of the architecture described in this patent. The following can be called an RDMA capability or an RDMA mechanism or an RDMA function. In such a system the application running on the initiator or target registers a region of memory, block 702, which is made available to its peer(s) for access directly from the NIC/HBA without substantial host intervention. These applications would also let their

peer(s) know about the memory regions being available for RDMA, block 708. Once both peers of the communication are ready to use the RDMA mechanism, the data transfer from RDMA regions can happen with essentially zero copy overhead from the source to the destination without substantial host intervention if NIC/HBA hardware in the peers implement

5 RDMA capability. The source, or initiator, would inform its peer of its desire to read or write specific RDMA enabled buffers and then let the destination or target, push or pull the data to/from its RDMA buffers. The initiator and the target NIC/HBA would then transport the data using the TCP/IP hardware implementation described in this patent, RMDA 703, TCP/IP offload 704, RMDA 708 and TCP/IP offload 709, between each other without substantial

10 intervention of the host processors, thereby significantly reducing the processor overhead. This mechanism would significantly reduce the TCP/IP processing overhead on the host processor and eliminate the need for multiple buffer copies for the data transfer illustrated in Fig. 6. RDMA enabled systems would thus allow the system, whether fast or slow, to perform the data transfer without creating a performance bottleneck for its peer. RDMA

15 capability implemented in this processor in storage over IP solution eliminates host intervention except usually at the data transfer start and termination. This relieves the host processors in both target and initiator systems to perform useful tasks without being interrupted at each packet arrival or transfer. RDMA implementation also allows the system to be secure and prevent unauthorized access. This is accomplished by registering the

20 exported memory regions with the HBA/NIC with their access control keys along with the region IDs. The HBA/NIC performs the address translation of the memory region request from the remote host to the RDMA buffer, performs security operations such as security key verification and then allows the data transfer. This processing is performed off the host processor in the processor of this invention residing on the HBA/NIC or as a companion

25 processor to the host processor on the motherboard, for example. This capability can also be used for large data transfers for server clustering applications as well as client server applications. Real time media applications transferring large amounts of data between a source or initiator and a destination or target can benefit from this.

Fig. 8 illustrates the host file system and SCSI stack implemented in software. As indicated

30 earlier the IP based storage stack, blocks 805, 806, 807, 808 and 809, should represent a consistent interface to the SCSI layers, blocks 803 and 804, as that provided by SCSI transport layer, block 811, or Fibre channel transport, block 810. This figure illustrates high level requirements that are imposed on the IP based storage implementation from a system level, besides those imposed by various issues of IP which is not designed to transport

35 performance sensitive block data.

Fig. 9 illustrates the iSCSI stack in more detail from that illustrated in Fig. 8. The iSCSI stack blocks 805 through 809, should provide an OS defined driver interface level functionality to the SCSI command consolidation layer blocks 803 & 804, such that the behavior of this layer and other layers on top of it are unchanged. Fig. 9 illustrates a set of functions that would be implemented to provide IP storage capabilities. The functions that provide the iSCSI functionality are grouped into related sets of functions, although there can be many variations of these as any person skilled in this area would appreciate. There are a set of functions that are required to meet the standard (e.g. target and initiator login and logout) functions, block 916, connection establishment and teardown functions, block 905. The figure illustrates functions that allow the OS SCSI software stack to discover the iSCSI device, block 916, set and get options/parameters, blocks 903 and 909, to start the device, block 913 and release the device, block 911. Besides the control functions discussed earlier, the iSCSI implementation provides bulk data transfer functions, through queues 912 and 917, to transport the PDUs specified by the iSCSI standard. The iSCSI stack may also include direct data transfer/placement (DDT) or RDMA functions or combination thereof, block 918, which are used by the initiator and target systems to perform substantially zero buffer copy and host intervention-less data transfers including storage and other bulk block data transfers. The SCSI commands and the block data transfers related to these are implemented as command queues, blocks 912 and 917, which get executed on the described processor. The host is interrupted primarily on the command completion. The completed commands are queued for the host to act on at a time convenient to the host. The figure illustrates the iSCSI protocol layer and the driver layer layered on the TCP/IP stack, blocks 907 and 908, which is also implemented off the host processor on the IP processor system described herein.

Fig. 10 illustrates the TCP/IP stack functionality that is implemented in the described IP processor system. These functions provide an interface to the upper layer protocol functions to carry the IP storage traffic as well as other applications that can benefit from direct OS TCP/IP bypass, RDMA or network sockets direct capabilities or combination thereof to utilize the high performance TCP/IP implementation of this processor. The TCP/IP stack provides capabilities to send and receive upper layer data, blocks 1017 and 1031, and command PDUs, establish the transport connections and teardown functions, block 1021, send and receive data transfer functions, checksum functions, block 1019, as well as error handling functions, block 1022, and segmenting and sequencing and windowing operations, block 1023. Certain functions like checksum verification/creation touch every byte of the data transfer whereas some functions that transport the data packets and update the transmission control block or session data base are invoked for each packet of the data

transfer. The session DB, block 1025, is used to maintain various information regarding the active sessions/connections along with the TCP/IP state information. The TCP layer is built on top of IP layer that provides the IP functionality as required by the standard. This layer provides functions to fragment/de-fragment, block 1033, the packets as per the path MTU, providing the route and forwarding information, block 1032, as well as interface to other functions necessary for communicating errors like, for example, ICMP, block 1029. The IP layer interfaces with the Ethernet layer or other media access layer technology to transport the TCP/IP packets onto the network. The lower layer is illustrated as Ethernet in various figures in this description, but could be other technologies like SONET, for instance, to transport the packets over SONET on MANs/WANs. Ethernet may also be used in similar applications, but may be used more so within a LAN and dedicated local SAN environments, for example.

Fig. 11 illustrates the iSCSI data flow. The figure illustrates the receive and transmit path of the data flow. The Host's SCSI command layer working with the iSCSI driver, both depicted in block 1101, would schedule the commands to be processed to the command scheduler, block 1108, in the storage flow controller seen in more detail in Fig. 26. The command scheduler 1108 schedules the new commands for operation in the processor described in more detail in Fig. 17. A new command that is meant for the target device with an existing connection gets en-queued to that existing connection, block 1111. When the connection to the target device does not exist, a new command is en-queued on to the unassigned command queue, block 1102. The session/connection establishment process like that shown in Fig. 47 and blocks 905 and 1006 is then called to connect to the target. Once the connection is established the corresponding command from the queue 1102 gets en-queued to the newly created connection command queue 1111 by the command scheduler 1108 as illustrated in the figure. Once a command reaches a stage of execution, the receive 1107 or transmit 1109 path is activated depending on whether the command is a read or a write transaction. The state of the connection/session which the command is transported is used to record the progress of the command execution in the session database as described subsequently. The buffers associated with the data transfer may be locked till such time as the transfer is completed. If the RDMA mechanism is used to transfer the data between the initiator and the target, appropriate region buffers identifiers, access control keys and related RDMA state data is maintained in memory on board the processor and may also be maintained in off-chip memory depending on the implementation chosen. As the data transfer, which may be over multiple TCP segments, associated with the command is completed the status of the command execution is passed onto the host SCSI layer which then does the appropriate processing. This may involve releasing the buffers being used for

data transfers to the applications, statistics update, and the like. During transfer, the iSCSI PDUs are transmitted by the transmit engines, block 1109, working with the transmit command engines, block 1110, that interpret the PDU and perform appropriate operations like retrieving the application buffers from the host memory using DMA to the storage processor and keeping the storage command flow information in the iSCSI connection database updated with the progress. As used in this patent the term "engine" can be a data processor or a part of a data processor, appropriate for the function or use of the engine. Similarly, the receive engines, block 1107, interpret the received command into new requests, response, errors or other command or data PDUs that need to be acted on appropriately. These receive engines working with the command engines, block 1106, route the read data or received data to the appropriate allocated application buffer through direct data transfer/placement or RDMA control information maintained for the session in the iSCSI session table. On command completion the control to the respective buffers, blocks 1103 and 1112, is released for the application to use. Receive and transmit engines can be the SAN packet processors 1706(a) to 1706(n) of Fig. 17 of this IP processor working with the session information recorded in the session data base entries 1704, which can be viewed as a global memory as viewed from the TCP/IP processor of Fig. 23 or the IP processor of Fig. 24. The same engines can get reused for different packets and commands with the appropriate storage flow context provided by the session database discussed in more detail below with respect to block 1704 and portion of session database in 1708 of Fig. 17. For clarification, the terms IP network application processor, IP Storage processor, IP Storage network application processor and IP processor can be the same entity, depending on the application. An IP network application processor core or an IP storage network application processor core can be the same entity, depending on the application.

Similarly a control command can use the transmit path whereas the received response would use the receive path. Similar engines can exist on the initiator as well as the target. The data flow direction is different depending on whether it is the initiator or the target. However, primarily similar data flow exists on both initiator and target with additional steps at the target. The target needs to perform additional operations to reserve the buffers needed to get the data of a write command, for instance, or may need to prepare the read data before the data is provided to the initiator. Similar instances would exist in case of an intermediate device, although, in such a device, which may be a switch or an appliance, some level of virtualization or frame filtering or such other operation may be performed that may require termination of the session on one side and originating sessions on the other.

This functionality is supported by this architecture but not illustrated explicitly in this figure, inasmuch as it is well within the knowledge of one of ordinary skill in the art.

Fig. 12 through Fig. 15 illustrate certain protocol information regarding transport sessions and how that information may be stored in a database in memory.

Fig. 12 illustrates the data structures that are maintained for iSCSI protocol and associated TCP/IP connections. The data belonging to each iSCSI session, block 1201, which is
5 essentially a nexus of initiator and target connections, is carried on the appropriate connection, block 1202. Dependent commands are scheduled on the queues of the same connection to maintain the ordering of the commands, block 1203. However, unrelated commands can be assigned to different transport connection. It is possible to have all the commands be queued to the same connection, if the implementation supports only one
10 connection per session. However, multiple connections per session are feasible to support line trunking between the initiator and the target. For example, in some applications, the initiator and the target will be in communication with each other and will decide through negotiation to accept multiple connections. In others, the initiator and target will communicate through only one session or connection. Fig. 13 and Fig. 14 illustrate the
15 TCP/IP and iSCSI session data base or transmission control block per session and connection. These entries may be carried as separate tables or may be carried together as a composite table as seen subsequently with respect to Figs. 23, 24, 26 and 29 depending on the implementation chosen and the functionality implemented e.g. TCP/IP only, TCP/IP with RDMA, IP Storage only, IP storage with TCP/IP, IP Storage with RDMA and the like.
20 Various engines that perform TCP/IP and storage flow control use all or some of these fields or more fields not shown, to direct the block data transfer over TCP/IP. The appropriate fields are updated as the connection progresses through the multiple states during the course of data transfer. Fig. 15 illustrates one method of storing the transmission control entries in a memory subsystem that consists of an on-chip session cache, blocks 1501 and
25 1502, and off-chip session memory, blocks 1503, 1504, 1505, 1506 and 1507, that retains the state information necessary for continuous progress of the data transfers.

Fig. 16 illustrates the IP processor architecture at a high level of abstraction. The processor consists of modular and scalable IP network application processor core, block 1603. Its functional blocks provide the functionality for enabling high speed storage and data transport
30 over IP networks. The processor core can include an intelligent flow controller, a programmable classification engine and a storage/network policy engine. Each can be considered an individual processor or any combination of them can be implemented as a single processor. The disclosed processor also includes a security processing block to provide high line rate encryption and decryption functionality for the network packets. This,
35 likewise, can be a single processor, or combined with the others mentioned above. The

disclosed processor includes a memory subsystem, including a memory controller interface, which manages the on chip session cache/memory, and a memory controller, block 1602, which manages accesses to the off chip memory which may be SRAM, DRAM, FLASH, ROM, EEPROM, DDR SDRAM, RDRAM, FCRAM, QDR SRAM, or other derivatives of static or dynamic random access memory or a combination thereof. The IP processor includes appropriate system interfaces to allow it to be used in the targeted market segments, providing the right media interfaces, block 1601, for LAN, SAN, WAN and MAN networks, and similar networks, and appropriate host interface, block 1606. The media interface block and the host interface block may be in a multi-port form where some of the ports may serve the redundancy and fail-over functions in the networks and systems in which the disclosed processor is used. The processor also may contain the coprocessor interface block 1605, for extending the capabilities of the main processor for example creating a multi-processor system. The system controller interface of block 1604 allows this processor to interface with an off-the-shelf microcontroller that can act as the system controller for the system in which the disclosed processor may be used. The processor architecture also support a control plane processor on board, that could act as the system controller or session manager. The system controller interface may still be provided to enable the use of an external processor. Such a version of this processor may not include the control processor for die cost reasons. There are various types of the core architecture that can be created, targeting specific system requirements, for example server adapters or storage controllers or switch line cards or other networking systems. The primary differences would be as discussed in the earlier sections of this patent. These processor blocks provide capabilities and performance to achieve the high performance IP based storage using standard protocols like iSCSI, FCIP, iFCP and the like. The detailed architecture of these blocks will be discussed in the following description.

Fig. 17 illustrates the IP processor architecture in more detail. The architecture provides capabilities to process incoming IP packets from the media access control (MAC) layer, or other appropriate layer, through full TCP/IP termination and deep packet inspection. This block diagram does not show the MAC layer block 1601, or blocks 1602, 1604 or 1605 of Fig. 16. The MAC layer interface blocks to the input queue, block 1701, and output queue, block 1712, of the processor in the media interface, block 1601, shown in Fig. 16. The MAC functionality could be standards based, with the specific type dependent on the network. Ethernet and Packet over SONET are examples of the most widely used interfaces today which may be included on the same silicon or a different version of the processor created with each.

The block diagram in Fig. 17 illustrates input queue and output queue blocks 1701 and 1712 as two separate blocks. The functionality may be provided using a combined block. The input queue block 1701 consists of the logic, control and storage to retrieve the incoming packets from the MAC interface block. Block 1701 queues the packets as they arrive from the interface and creates appropriate markers to identify start of the packet, end of the packet and other attributes like a fragmented packet or a secure packet, and the like, working with the packet scheduler 1702 and the classification engine 1703. The packet scheduler 1702, can retrieve the packets from the input queue controller and passes them for classification to the classification engine. The classification block 1703, is shown to follow the scheduler, however from a logical perspective the classification engine receives the packet from the input queue, classifies the packet and provides the classification tag to the packet, which is then scheduled by the scheduler to the processor array 1706(a) . . . 1706(n). Thus the classification engine can act as a pass-through classification engine, sustaining the flow of the packets through its structure at the full line rate. The classification engine is a programmable engine that classifies the packets received from the network in various categories and tags the packet with the classification result for the scheduler and the other packet processors to use. Classification of the network traffic is a very compute intensive activity which can take up to half of the processor cycles available in a packet processor. This integrated classification engine is programmable to perform Layer 2 through Layer 7 inspection. The fields to be classified are programmed in with expected values for comparison and the action associated with them if there is a match. The classifier collects the classification walk results and can present these as a tag to the packet identifying the classification result as seen subsequently with respect to Fig. 30. This is much like a tree structure and is understood as a "walk." The classified packets are then provided to the scheduler 1702 as the next phase of the processing pipeline.

The packet scheduler block 1702 includes a state controller and sequencer that assign packets to appropriate execution engines on the disclosed processor. The execution engines are the SAN packet processors, block 1706(a) through 1706(n), including the TCP/IP and/or storage engines as well as the storage flow/RDMA controller, block 1708 or host bypass and/or other appropriate processors, depend on the desired implementation. For clarity, the term "/", when used to designate hardware components in this patent, can mean "and/or" as appropriate. For example, the component "storage flow/RDMA controller" can be a storage flow and RDMA controller, a storage flow controller, or an RDMA controller, as appropriate for the implementation. The scheduler 1702 also maintains the packet order through the processor where the state dependency from a packet to a packet on the same connection/session is important for correct processing of the incoming packets. The

scheduler maintains various tables to track the progress of the scheduled packets through the processor until packet retirement. The scheduler also receives commands that need to be scheduled to the packet processors on the outgoing commands and packets from the host processor or switch fabric controller or interface.

5 The TCP/IP and storage engines along with programmable packet processors are together labeled as the SAN Packet Processors 1706(a) through 1706(n) in Fig. 17. These packet processors are engines that are independent programmable entities that serve a specific role. Alternatively, two or more of them can be implemented as a single processor depending on the desired implementation. The TCP/IP engine of Fig. 23 and the storage
10 engines of Fig. 24 are configured in this example as coprocessors to the programmable packet processor engine block 2101 of Fig. 21. This architecture can thus be applied with relative ease to applications other than storage by substituting/removing for the storage engine for reasons of cost, manufacturability, market segment and the like. In a pure networking environment the storage engine could be removed, leaving the packet processor
15 with a dedicated TCP/IP engine and be applied for the networking traffic, which will face the same processing overhead from TCP/IP software stacks. Alternatively one or more of the engines may be dropped for desired implementation e.g. for processor supporting only IP Storage functions may drop TCP/IP engine and/or packet engine which may be in a separate chip. Hence, multiple variations of the core scalable and modular architecture are
20 possible. The core architecture can thus be leveraged in applications beside the storage over IP applications by substituting the storage engine with other dedicated engines, for example a high performance network security and policy engine, a high performance routing engine, a high performance network management engine, deep packet inspection engine providing string search, an engine for XML, an engine for virtualization, and the like,
25 providing support for an application specific acceleration. The processing capability of this IP processor can be scaled by scaling the number of SAN Packet Processor blocks 1706 (a) through 1706 (n) in the chip to meet the line rate requirements of the network interface. The primary limitation from the scalability would come from the silicon real-estate required and the limits imposed by the silicon process technologies. Fundamentally this architecture is
30 scalable to very high line rates by adding more SAN packet processor blocks thereby increasing the processing capability. Other means of achieving a similar result is to increase the clock frequency of operation of the processor to that feasible within the process technology limits.

Fig. 17 also illustrates the IP session cache/memory and the memory controller block 1704.
35 This cache can be viewed as an internal memory or local session database cache. This

block is used to cache and store the TCP/IP session database and also the storage session database for a certain number of active sessions. The number of sessions that can be cached is a direct result of the chosen silicon real-estate and what is economically feasible to manufacture. The sessions that are not on chip, are stored and retrieved to/from off chip memory, viewed as an external memory, using a high performance memory controller block which can be part of block 1704 or otherwise. Various processing elements of this processor share this controller using a high speed internal bus to store and retrieve the session information. The memory controller can also be used to temporarily store packets that may be fragmented or when the host interface or outbound queues are backed-up. The controller may also be used to store statistics information or any other information that may be collected by the disclosed processor or the applications running on the disclosed or host processor.

The processor block diagram of Fig. 17 also illustrates host interface block 1710, host input queue, block 1707 and host output queue, block 1709 as well as the storage flow / RDMA controller, block 1708. These blocks provide the functions that are required to transfer data to and from the host (also called "peer") memory or switch fabric. These blocks also provide features that allow the host based drivers to schedule the commands, retrieve incoming status, retrieve the session database entry, program the disclosed processor, and the like to enable capabilities like sockets direct architecture, full TCP/IP termination, IP storage offload and the like capabilities with or without using RDMA. The host interface controller 1710, seen in greater detail in Fig. 27, provides the configuration registers, DMA engines for direct memory to memory data transfer, the host command block that performs some of the above tasks, along with the host interface transaction controller and the host interrupt controller. The host input and output queues 1707, 1709 provide the queuing for incoming and outgoing packets. The storage flow and RDMA controller block 1708 provides the functionality necessary for the host to queue the commands to the disclosed processor, which then takes these commands and executes them, interrupting the host processor on command termination. The RDMA controller portion of block 1708 provides various capabilities necessary for enabling remote direct memory access. It has tables that include information such as RDMA region, access keys, and virtual address translation functionality. The RDMA engine inside this block performs the data transfer and interprets the received RDMA commands to perform the transaction if the transaction is allowed. The storage flow controller of block 1708 also keeps track of the state of the progress of various commands that have been scheduled as the data transfer happens between the target and the initiator. The storage flow controller schedules the commands for execution and also provides the command completion information to the host drivers. The above can be considered RDMA

capability and can be implemented as described or by implementing as individual processors, depending on designer's choice. Also, additional functions can be added to or removed from those described without departing from the spirit or the scope of this patent.

5 The control plane processor block 1711 of this processor is used to provide relatively slow path functionality for TCP/IP and/or storage protocols which may include error processing with ICMP protocol, name resolution, address resolution protocol, and it may also be programmed to perform session initiation/teardown acting as a session controller/connection manager, login and parameter exchange, and the like. This control plane processor could be off chip to provide the system developer a choice of the control plane processor, or may be
10 on chip to provide an integrated solution. If the control plane processor is off-chip, then an interface block would be created or integrated herein that would allow this processor to interface with the control plane processor and perform data and command transfers. The internal bus structures and functional block interconnections may be different than illustrated for all the detailed figures for performance, die cost requirements and the like and not depart
15 from the spirit and the scope of this patent.

Capabilities described above for Fig. 17 blocks with more detail below, enable a packet streaming architecture that allows packets to pass through from input to output with minimal latency, with in-stream processing by various processing resources of the disclosed processor.

20 Fig. 18 illustrates the input queue and controller block shown generally at 1701 of Fig. 17 in more detail. The core functionality of this block is to accept the incoming packets from multiple input ports, Ports 1 to N, in blocks 1801 and 1802(i) to 1802(n), and to queue them using a fixed or programmable priority on the input packet queue, block 1810, from where the packets get de-queued for classifier, scheduler and further packet processing through
25 scheduler I/F blocks 1807-1814. The input queue controller interfaces with each of the input ports (Port 1 through Port N in a multi-port implementation), and queues the packets to the input packet queue 1810. The packet en-queue controller and marker block 1804 may provide fixed priority functions or may be programmable to allow different policies to be applied to different interfaces based on various characteristics like port speed, the network
30 interface of the port, the port priority and others that may be appropriate. Various modes of priority may be programmable like round-robin, weighted round-robin or others. The input packet de-queue controller 1812 de-queues the packets and provides them to the packet scheduler, block 1702 of Fig. 17 via scheduler I/F 1814. The scheduler schedules the packets to the SAN packet processors 1706 (a) – 1706 (n) once the packets have been
35 classified by the classification engine 1703 of Fig. 17. The encrypted packets can be

classified as encrypted first and passed on to the security engine 1705 of Fig. 17 by the secure packet interface block 1813 of Fig. 18. for authentication and/or decryption if the implementation includes security processing otherwise the security interfaces may not be present and an external security processor would be used to perform similar functions. The

5 decrypted packets from clear packet interface, block 1811, are then provided to the input queue through block 1812 from which the packet follows the same route as a clear packet. The fragmented IP packets may be stored on-chip in the fragmented packet store and controller buffers, block 1806, or may be stored in the internal or external memory. When the last fragment arrives, the fragment controller of block 1806, working with the

10 classification engine and the scheduler of Fig. 17, merges these fragments to assemble the complete packet. Once the fragmented packet is combined to form a complete packet, the packet is scheduled into the input packet queue via block 1804 and is then processed by the packet de-queue controller, block 1812, to be passed on to various other processing stages of this processor. The input queue controller of Fig. 18 assigns a packet tag/descriptor to

15 each incoming packet which is managed by the attribute manager of block 1809 which uses the packet descriptor fields like the packet start, size, buffer address, along with any other security information from classification engine, and stored in the packet attributes and tag array of block 1808. The packet tag and attributes are used to control the flow of the packet through the processor by the scheduler and other elements of the processor in an efficient

20 manner through interfaces 1807, 1811, 1813 and 1814

Fig. 19 illustrates the packet scheduler and sequencer 1702 of Fig. 17 in more detail. This block is responsible for scheduling packets and tasks to the execution resources of this processor and thus also acts as a load balancer. The scheduler retrieves the packet headers from the header queue, block 1902, from the input queue controller 1901 to pass

25 them to the classification engine 1703 of Feb. 17 which returns the classification results to the classifier queue, block 1909, that are then used by the rest of the processor engines. The classification engine may be presented primarily with the headers, but if deep packet inspection is also programmed, the classification engine may receive the complete packets which it routes to the scheduler after classification. The scheduler comprises a classification

30 controller/scheduler, block 1908, which manages the execution of the packets through the classification engine. This block 1908 of Fig. 19 provides the commands to the input queue controller, block 1901, in case of fragmented packets or secure packets, to perform the appropriate actions for such packets e.g. schedule an encrypted packet to the security engine of Fig. 17. The scheduler state control and the sequencer, block 1916, receive state

35 information of various transactions/operations active inside the processor and provide instructions for the next set of operations. For instance, the scheduler retrieves the packets

from the input packet queue of block 1903, and schedules these packets in the appropriate resource queue depending on the results of the classification received from the classifier or directs the packet to the packet memory, block 1913 or 1704 through 1906, creating a packet descriptor/tag which may be used to retrieve the packet when appropriate resource

5 needs it to performs its operations at or after scheduling. The state control and sequencer block 1916 instructs/directs the packets with their classification result, block 1914, to be stored in the packet memory, block 1913, from where the packets get retrieved when they are scheduled for operation. The state controller and the sequencer identify the execution resource that should receive the packet for operation and creates a command and assigns

10 this command with the packet tag to the resource queues, blocks 1917 (Control Plane), 1918 (port i-port n), 1919 (bypass) and 1920 (host) of Fig. 19. The priority selector 1921 is a programmable block that retrieves the commands and the packet tag from the respective queues based on the assigned priority and passes this to the packet fetch and command controller, block 1922. This block retrieves the packet from the packet memory store 1913

15 along with the classification results and schedules the packet transfer to the appropriate resource on the high performance processor command and packet busses such as at 1926 when the resource is ready for operation. The bus interface blocks, like command bus interface controller 1905, of the respective recipients interpret the command and accept the packet and the classification tag for operation. These execution engines inform the

20 scheduler when the packet operation is complete and when the packet is scheduled for its end destination (either the host bus interface, or the output interface or control plane interface, etc.). This allows the scheduler to retire the packet from its state with the help of retirement engine of block 1904 and frees up the resource entry for this session in the resource allocation table, block 1923. The resource allocation table is used by the

25 sequencer to assign the received packets to specific resources, depending on the current state of internal state of these resources, e.g. the session database cache entry buffered in the SAN packet processor engine, the connection ID of the current packet being executed in the resource, and the like. Thus packets that are dependent on an ordered execution get assigned primarily to the same resource, which improves memory traffic and performance by

30 using the current DB state in the session memory in the processor and not have to retrieve new session entries. The sequencer also has interface to the memory controller, block 1906, for queuing of packets that are fragmented packets and/or for the case in which the scheduler queues get backed-up due to a packet processing bottleneck down stream, which may be caused by specific applications that are executed on packets that take more time

35 than that allocated to maintain a full line rate performance, or for the case in which any other downstream systems get full, unable to sustain the line rate.

If the classifier is implemented before the scheduler as discussed above with respect to Fig. 17 where the classification engine receives the packet from the input queue, items 1901, 1902, 1908, 1909 and 1910 would be in the classifier, or may not be needed, depending on the particular design. The appropriate coupling from the classifier to/from the scheduler
5 blocks 1903, 1907, 1914 and 1915 may be created in such a scenario and the classifier coupled directly to the input queue block of Fig. 18.

Fig. 20 illustrates the packet classification engine shown generally at 1703 of Fig. 17. Classification of the packets into their various attributes is a very compute intensive operation. The classifier can be a programmable processor that examines various fields of
10 the received packet to identify the type of the packet, the protocol type e.g. IP, ICMP, TCP, UDP etc, the port addresses, the source and destination fields, etc. The classifier can be used to test a particular field or a set of fields in the header or the payload. The block diagram illustrates a content addressable memory based classifier. However, as discussed earlier this could be a programmable processor as well. The primary differences are the
15 performance and complexity of implementation of the engine. The classifier gets the input packets through the scheduler from the input queues, blocks 2005 and 2004 of Fig. 20. The input buffers 2004 queue the packets/descriptor and/or the packet headers that need to be classified. Then the classification sequencer 2003 fetches the next available packet in the queue and extracts the appropriate packet fields based on the global field descriptor sets,
20 block 2007, which are, or can be, programmed. Then the classifier passes these fields to the content addressable memory (CAM) array, block 2009, to perform the classification. As the fields are passed through the CAM array, the match of these fields identifies next set of fields to be compared and potentially their bit field location. The match in the CAM array results in the action/event tag, which is collected by the result compiler, (where "compiling" is
25 used in the sense of "collecting") block 2014 and also acted on as an action that may require updating the data in the memory array, block 2013, associated with specific CAM condition or rule match. This may include performing an arithmetic logic unit (ALU) operation, block 2017, which can be considered one example of an execution resource) on this field e.g. increment or decrement the condition match and the like. The CAM arrays are programmed
30 with the fields, their expected values and the action on match, including next field to compare, through the database initialization block 2011, accessible for programming through the host or the control plane processor interfaces 1710, 1711. Once the classification reaches a leaf node the classification is complete and the classification tag is generated that identifies the path traversed that can then be used by other engines of the IP processor
35 avoid performing the same classification tasks. For example a classification tag may include the flow or session ID, protocol type indication e.g. TCP/UDP/ICMP etc., value indicating

whether to processes, bypass, drop packet, drop session, and the like, or may also include the specific firmware code routine pointer for the execution resource to start packet processing or may include signature of the classification path traversed or the like. The classification tag fields are chosen based on processor implementation and functionality.

5 The classifier retirement queue, block 2015, holds the packets/descriptors of packets that are classified and classification tag and are waiting to be retrieved by the scheduler. The classification data base can be extended using database extension interface and pipeline control logic block 2006. This allows systems that need extensibility for a larger classification database to be built. The classification engine with the action interpreter, the

10 ALU and range matching block of 2012 also provide capabilities to program storage / network policies / actions that need to be taken if certain policies are met. The policies can be implemented in the form of rule and action tables. The policies get compiled and programmed in the classification engine through the host interface along with the classification tables. The database interface and pipeline control 2006 could be

15 implemented to couple to companion processor to extend the size of the classification/policy engine.

Fig. 21 illustrates the SAN Packet Processor shown generally at 1706 (a) through 1706 (n) of Fig. 17. A packet processor can be a specially designed packet processor, or it can be any suitable processor such as an ARM, ARC, Tensilica, MIPS, StrongARM, X86, PowerPC,

20 Pentium processor, iA64 or any other processor that serves the functions described herein. This is also referred as the packet processor complex in various sections of this patent. This packet processor comprises a packet engine, block 2101, which is generally a RISC OR VLIW machine with target instructions for packet processing or a TCP/IP engine, block 2102 or an IP storage engine, block 2103 or a combination thereof. . These engines can be

25 configured as coprocessors to the packet engine or can be independent engines. Fig. 22 illustrates the packet engine in more detail. The packet engine is a generally RISC OR VLIW machine as indicated above with instruction memory, block 2202, and Data Memory, block 2206, (both of which can be RAM) that are used to hold the packet processing micro routines and the packets and intermediate storage. The instruction memory 2202 which, like

30 all such memory in this patent, can be RAM or other suitable storage, is initialized with the code that is executed during packet processing. The packet processing code is organized as tight micro routines that fit within the allocated memory. The instruction decoder and the sequencer, block 2204, fetches the instructions from instruction memory 2202, decodes them and sequences them through the execution blocks contained within the ALU, block

35 2208. This machine can be a simple pipelined engine or a more complex deep pipelined machine that may also be designed to provide a packet oriented instruction set. The DMA

engine, block 2205 and the bus controller, block 2201, allow the packet engine to move the data packets from the scheduler of Fig. 19 and the host interface into the data memory 2206 for operation. The DMA engine may hold multiple memory descriptors to store/retrieve packet/data to/from host memory/packet memory. This would enable memory accesses to happen in parallel to packet processor engine operations. The DMA engine 2205 also may be used to move the data packets to and from the TCP and storage engines 2210, 2211. Once the execution of the packet is complete, the extracted data or newly generated packet is transferred to the output interface either towards the media interface or the host interface

Fig. 23 illustrates a programmable TCP/IP packet processor engine, seen generally at 2210 of Fig. 22, in more detail. This engine is generally a programmable processor with common RISC OR VLIW instructions along with various TCP/IP oriented instructions and execution engines but could also be a micro-coded or a state machine driven processor with appropriate execution engines described in this patent. The TCP processor includes a checksum block, 2311, for TCP checksum verification and new checksum generation by executing these instructions on the processor. The checksum block extracts the data packet from the packet buffer memory (a Data RAM is one example of such memory), 2309, and performs the checksum generation or verification. The packet look-up interface block, 2310, assists the execution engines and the instruction sequencer, 2305, providing access to various data packet fields or the full data packet. The classification tag interpreter, 2313, is used by the instruction decoder 2304 to direct the program flow based on the results of the classification if such an implementation is chosen. The processor provides specific sequence and windowing operations including segmentation, block 2315, for use in the TCP/IP data sequencing calculations for example, to look-up the next expected sequence number and see if that received is within the agreed upon sliding window, which sliding window is a well known part of the TCP protocol, for the connection to which the packet belongs. This element 2315 may also include a segmentation controller like that shown at 2413 of Fig. 24. Alternatively, one of ordinary skill in the art, with the teaching of this patent, can easily implement the segmentation controllers elsewhere on the TCP/IP processor of this Fig. 23. The processor provides a hash engine, block 2317, which is used to perform hash operations against specific fields of the packet to perform a hash table walk that may be required to get the right session entry for the packet. The processor also includes a register file, block 2316, which extracts various commonly used header fields for TCP processing, along with pointer registers for data source and destination, context register sets, and registers that hold the TCP states along with a general purpose register file. The TCP/IP processor can have multiple contexts for packet execution, so that when a given packet execution stalls for any reason, for example memory access, the other context can

be woken up and the processor continue the execution of another packet stream with little efficiency loss. The TCP/IP processor engine also maintains a local session cache, block 2320, which holds most recently used or most frequently used entries, which can be used locally without needing to retrieve them from the global session memory. The local session

5 cache can be considered an internal memory of the TCP/IP processor, which can be a packet processor. Of course, the more entries that will be used that can be stored locally in the internal memory, without retrieving additional ones from the session, or global, memory, the more efficient the processing will be. The packet scheduler of Fig. 19 is informed of the connection IDs that are cached per TCP/IP processor resource, so that it can schedule the

10 packets that belong to the same session to the same packet processor complex. When the packet processor does not hold the session entry for the specific connection, then the TCP session database lookup engine, block 2319, working with the session manager, block 2321, and the hash engine retrieves the corresponding entry from the global session memory through the memory controller interface, block 2323. There are means, such as logic

15 circuitry inside the session manager that allow access of session entries or fields of session entries, that act with the hash engine to generate the session identifier for storing/retrieving the corresponding session entry or its fields to the session database cache. This can be used to update those fields or entries as a result of packet processing. When a new entry is fetched, the entry which it is replacing is stored to the global session memory. The local

20 session caches may follow exclusivity caching principles, so that multiple processor complexes do not cause any race conditions, damaging the state of the session. Other caching protocols like MESI protocol may also be used to achieve similar results. When a session entry is cached in a processor complex, and another processor complex needs that entry, this entry is transferred to the new processor with exclusive access or appropriate

25 caching state based on the algorithm. The session entry may also get written to the global session memory in certain cases. The TCP/IP processor also includes a TCP state machine, block 2322, which is used to walk through the TCP states for the connection being operated on. This state machine receives the state information stored in the session entry along with the appropriate fields affecting the state from the newly received packet. This

30 allows the state machine to generate the next state if there is a state transition and the information is updated in the session table entry. The TCP/IP processor also includes a frame controller/out of order manager block, 2318, that is used to extract the frame information and perform operations for out of order packet execution. This block could also include an RDMA mechanism such as that shown at 2417 of Fig. 24, but used for non-

35 storage data transfers. One of ordinary skill in the art can also, with the teaching of this patent, implement an RDMA mechanism elsewhere on the TCP/IP processor. This architecture creates an upper layer framing mechanism which may use packet CRC as

framing key or other keys that is used by the programmable frame controller to extract the embedded PDUs even when the packets arrive out of order and allow them to be directed to the end buffer destination. This unit interacts with the session database to handle out of order arrival information which is recorded so that once the intermediate segments arrive, the retransmissions are avoided. Once the packet has been processed through the TCP/IP processor, it is delivered for operation to the storage engine, if the packet belongs to a storage data transfer and the specific implementation includes a storage engine, otherwise the packet is passed on to the host processor interface or the storage flow/RDMA controller of block 1708 for processing and for DMA to the end buffer destination. The packet may be transferred to the packet processor block as well for any additional processing on the packet. This may include application and customer specific application code that can be executed on the packet before or after the processing by the TCP/IP processor and the storage processor. Data transfer from the host to the output media interface would also go through the TCP/IP processor to form the appropriate headers to be created around the data and also perform the appropriate data segmentation, working with the frame controller and/or the storage processor as well as to update the session state. This data may be retrieved as a result of host command or received network packet scheduled by the scheduler to the packet processor for operation. The internal bus structures and functional block interconnections may be different than illustrated for performance, die cost requirements and the like. For example, Host Controller Interface 2301, Scheduler Interface 2307 and Memory Controller Interface 2323 may be part of a bus controller that allows transfer of data packets or state information or commands, or a combination thereof, to or from a scheduler or storage flow/RDMA controller or host or session controller or other resources such as, without limitation, security processor, or media interface units, host interface, scheduler, classification processor, packet buffers or controller processor, or any combination of the foregoing.

Fig. 24 illustrates the IP storage processor engine of Fig. 22 in more detail. The storage engine is a programmable engine with an instruction set that is geared towards IP based storage along with, usually, a normal RISC OR VLIW-like packet processing instruction set. The IP storage processor engine contains block 2411, to perform CRC operations. This block allows CRC generation and verification. The incoming packet with IP storage is transferred from the TCP/IP engine through DMA, blocks 2402 and 2408, into the data memory (a data RAM is an example of such memory), block 2409. When the implementation does not include TCP/IP engine or packet processor engine or a combination thereof, the packet may be received from the scheduler directly for example. The TCP session database information related to the connection can be retrieved from the

local session cache as needed or can also be received with the packet from the TCP/IP engine. The storage PDU is provided to the PDU classifier engine, block 2418, which classifies the PDU into the appropriate command, which is then used to invoke the appropriate storage command execution engine, block 2412. The command execution can
5 be accomplished using the RISC OR VLIW, or equivalent, instruction set or using a dedicated hardware engine. The command execution engines perform the command received in the PDU. The received PDU may contain read command data, or R2T for a pending write command or other commands required by the IP storage protocol. These engines retrieve the write data from the host interface or direct the read data to the
10 destination buffer. The storage session database entry is cached, in what can be viewed as a local memory, block 2420, locally for the recent or frequent connections served by the processor. The command execution engines execute the commands and make the storage database entry updates working with the storage state machine, block 2422, and the session manager, block 2421. The connection ID is used to identify the session, and if the session is
15 not present in the cache, then it is retrieved from the global session memory 1704 of Fig. 17 by the storage session look-up engine, block 2419. For data transfer from the initiator to target, the processor uses the segmentation controller, block 2413, to segment the data units into segments as per various network constraints like path MTU and the like. The segmentation controller attempts to ensure that the outgoing PDUs are optimal size for the
20 connection. If the data transfer requested is larger than the maximum effective segment size, then the segmentation controller packs the data into multiple packets and works with the sequence manager, block 2415, to assign the sequence numbers appropriately. The segmentation controller 2413 may also be implemented within the TCP/IP processor of Fig. 23. That is, the segmentation controller may be part of the sequence/window operations
25 manager 2315 of Fig. 23 when this processor is used for TCP/IP operations and not storage operations. One of ordinary skill in the art can easily suggest alternate embodiments for including the segmentation controller in the TCP/IP processor using the teachings of this patent. The storage processor of Fig. 24 (or the TCP/IP processor of Fig. 23) can also include an RDMA engine that interprets the remote direct memory access instructions
30 received in the PDUs for storage or network data transfers that are implemented using this RDMA mechanism. In Fig. 24, for example, this is RDMA engine 2417. In the TCP/IP processor of Fig. 23 an RDMA engine could be part of the frame controller and out of order manager 2318, or other suitable component. If both ends of the connection agree to the RDMA mode of data transfer, then the RDMA engine is utilized to schedule the data
35 transfers between the target and initiator without substantial host intervention. The RDMA transfer state is maintained in a session database entry. This block creates the RDMA headers to be layered around the data, and is also used to extract these headers from the

received packets that are received on RDMA enabled connections. The RDMA engine works with the storage flow/ RDMA controller, 1708, and the host interface controller, 1710, by passing the messages/instructions and performs the large block data transfers without substantial host intervention. The RDMA engine of the storage flow/RDMA controller block, 5 1708, of the IP processor performs protection checks for the operations requested and also provides conversion from the RDMA region identifiers to the physical or virtual address in the host space. This functionality may also be provided by RDMA engine, block 2417, of the storage engine of the SAN packet processor based on the implementation chosen. The distribution of the RDMA capability between 2417 and 1708 and other similar engines is an 10 implementation choice that one with ordinary skill in the art will be able to do with the teachings of this patent. Outgoing data is packaged into standards based PDU by the PDU creator, block 2425. The PDU formatting may also be accomplished by using the packet processing instructions. The storage engine of Fig. 24 works with the TCP/IP engine of Fig. 23 and the packet processor engine of Fig. 17 to perform the IP storage operations 15 involving data and command transfers in both directions i.e. from the initiator to target and the target to the host and vice versa. That is, the Host controller Interface 2401, 2407 store and retrieve commands or data or a combination thereof to or from the host processor. These interfaces may be directly connected to the host or may be connected through an intermediate connection. Though shown as two apparatus, interfaces 2401 and 2407 could 20 be implemented as a single apparatus. The flow of data through these blocks would be different based on the direction of the transfer. For instance, when command or data is being sent from the host to the target, the storage processing engines will be invoked first to format the PDU and then this PDU is passed on to the TCP processor to package the PDU in a valid TCP/IP segment. However, a received packet will go through the TCP/IP engine 25 before being scheduled for the storage processor engine. The internal bus structures and functional block interconnections may be different than illustrated for performance, die cost requirements, and the like. For example, and similarly to Fig. 23, Host Controller Interface 2401, 2407 and Memory Controller Interface 2423 may be part of a bus controller that allows transfer of data packets or state information or commands, or a combination thereof, to or 30 from a scheduler or host or storage flow/RDMA controller or session controller or other resources such as, without limitation, security processor, or media interface units, host interface, scheduler, classification processor, packet buffers or controller processor, or any combination of the foregoing.

In applications in which storage is done on a chip not including the TCP/IP processor of 35 Fig. 23 by, as one example, an IP Storage processor such as an iSCSI processor of Fig. 24, the TCP/IP Interface 2406 would function as an interface to a scheduler for scheduling IP

storage packet processing by the IP Storage processor. Similar variations are well within the knowledge of one of ordinary skill in the art, viewing the disclosure of this patent.

Fig. 25 illustrates the output queue controller block 1712 of Fig. 17 in more detail. This block receives the packets that need to be sent on to the network media independent interface
5 1601 of Fig. 16. The packets may be tagged to indicate if they need to be encrypted before being sent out. The controller queues the packets that need to be secured to the security engine through the queue 2511 and security engine interface 2510. The encrypted packets are received from the security engine and are queued in block 2509, to be sent to their destination. The output queue controller may assign packets onto their respective quality of
10 service (QOS) queues, if such a mechanism is supported. The programmable packet priority selector, block 2504, selects the next packet to be sent and schedules the packet for the appropriate port, Port1 . . . PortN. The media controller block 1601 associated with the port accepts the packets and sends them to their destination.

Fig. 26 illustrates the storage flow controller /RDMA controller block, shown generally at
15 1708 of Fig. 17, in more detail. The storage flow and RDMA controller block provides the functionality necessary for the host to queue the commands (storage or RDMA or sockets direct or a combination thereof) to this processor, which then takes these commands and executes them, interrupting the host processor primarily on command termination. The command queues, new and active, blocks 2611 and 2610, and completion queue, block
20 2612, can be partially on chip and partially in a host memory region or memory associated with the IP processor, from which the commands are fetched or the completion status deposited. The RDMA engine, block 2602, provides various capabilities necessary for enabling remote direct memory access. It has tables, like RDMA look-up table 2608, that include information like RDMA region and the access keys, and virtual address translation
25 functionality. The RDMA engine inside this block 2602 performs the data transfer and interprets the received RDMA commands to perform the transaction if allowed. The storage flow controller also keeps track of the state of the progress of various commands that have been scheduled as the data transfer happens between the target and the initiator. The storage flow controller schedules the commands for execution and also provides the
30 command completion information to the host drivers. The storage flow controller provides command queues where new requests from the host are deposited, as well as active commands are held in the active commands queue. The command scheduler of block 2601, assigns new commands, that are received which are for targets for which no connections exist, to the scheduler for initiating a new connection. The scheduler 1702, uses the control
35 plane processor shown generally at 1711 of Fig. 17 to do the connection establishment at

which point the connection entry is moved to the session cache, shown generally in Fig. 15 and 1704 in Fig. 17, and the state controller in the storage flow controller block 2601 moves the new command to active commands and associates the command to the appropriate connection. The active commands, in block 2610, are retrieved and sent to the scheduler, block 1702 for operation by the packet processors. The update to the command status is provided back to the flow controller which then stores it in the command state tables, blocks 2607 and accessed through block 2603. The sequencer of 2601 applies a programmable priority for command scheduling and thus selects the next command to be scheduled from the active commands and new commands. The flow controller also includes a new requests queue for incoming commands, block 2613. The new requests are transferred to the active command queue once the appropriate processing and buffer reservations are done on the host by the host driver. As the commands are being scheduled for execution, the state controller 2601 initiates data pre-fetch by host data pre-fetch manager, block 2617, from the host memory using the DMA engine of the host interface block 2707, hence keeping the data ready to be provided to the packet processor complex when the command is being executed. The output queue controller, block 2616, enables the data transfer, working with the host controller interface, block 2614. The storage flow/RDMA controller maintains a target-initiator table, block 2609, that associates the target/initiators that have been resolved and connections established for fast look-ups and for associating commands to active connections. The command sequencer may also work with the RDMA engine 2602, if the commands being executed are RDMA commands or if the storage transfers were negotiated to be done through the RDMA mechanism at the connection initiation. The RDMA engine 2602, as discussed above, provides functionality to accept multiple RDMA regions, access control keys and the virtual address translation pointers. The host application (which may be a user application or an OS kernel function, storage or non-storage such as downloading web pages, video files, or the like) registers a memory region that it wishes to use in RDMA transactions with the disclosed processor through the services provided by the associated host driver. Once this is done, the host application communicates this information to its peer on a remote end. Now, the remote machine or the host can execute RDMA commands, which are served by the RDMA blocks on both ends without requiring substantial host intervention. The RDMA transfers may include operations like read from a region, a certain number of bytes with a specific offset or a write with similar attributes. The RDMA mechanism may also include send functionality which would be useful in creating communication pipes between two end nodes. These features are useful in clustering applications where large amounts of data transfer is required between buffers of two applications running on servers in a cluster, or more likely, on servers in two different clusters of servers, or such other clustered systems. The storage data transfer may also be

accomplished using the RDMA mechanism, since it allows large blocks of data transfers without substantial host intervention. The hosts on both ends get initially involved to agree on doing the RDMA transfers and allocating memory regions and permissions through access control keys that get shared. Then the data transfer between the two nodes can
5 continue without host processor intervention, as long as the available buffer space and buffer transfer credits are maintained by the two end nodes. The storage data transfer protocols would run on top of RDMA, by agreeing to use RDMA protocol and enabling it on both ends. The storage flow controller and RDMA controller of Fig. 26 can then perform the storage command execution and the data transfer using RDMA commands. As the
10 expected data transfers are completed the storage command completion status is communicated to the host using the completion queue 2612. The incoming data packets arriving from the network are processed by the packet processor complex of Fig. 17 and then the PDU is extracted and presented to the flow controller OF FIG. 26 in case of storage/RDMA data packets. These are then assigned to the incoming queue block 2604,
15 and transferred to the end destination buffers by looking up the memory descriptors of the receiving buffers and then performing the DMA using the DMA engine inside the host interface block 2707. The RDMA commands may also go through protection key look-up and address translation as per the RDMA initialization.

The foregoing may also be considered a part of an RDMA capability or an RDMA
20 mechanism or an RDMA function.

Fig. 27 illustrates host interface controller 1710 of Fig. 17 in more detail. The host interface block includes a host bus interface controller, block 2709, which provides the physical interface to the host bus. The host interface block may be implemented as a fabric interface or media independent interface when embodied in a switch or a gateway or similar
25 configuration depending on the system architecture and may provide virtual output queuing and/or other quality of service features. The transaction controller portion of block 2708, executes various bus transactions and maintains their status and takes requested transactions to completion. The host command unit, block 2710, includes host bus configuration registers and one or more command interpreters to execute the commands
30 being delivered by the host. The host driver provides these commands to this processor over Host Output Queue Interface 2703. The commands serve various functions like setting up configuration registers, scheduling DMA transfers, setting up DMA regions and permissions if needed, setup session entries, retrieve session database, configure RDMA engines and the like. The storage and other commands may also be transferred using this
35 interface for execution by the IP processor.

Fig. 28 illustrates the security engine 1705 of Fig. 17 in more detail. The security engine illustrated provides authentication and encryption and decryption services like those required by standards like IPSEC for example. The services offered by the security engine may include multiple authentication and security algorithms. The security engine may be on-

5 board the processor or may be part of a separate silicon chip as indicated earlier. An external security engine providing IP security services would be situated in a similar position in the data flow, as one of the first stages of packet processing for incoming packets and as one of the last stages for the outgoing packet. The security engine illustrated provides advanced encryption standard (AES) based encryption and decryption services, which are

10 very hardware performance efficient algorithms adopted as security standards. This block could also provide other security capabilities like DES, 3DES, as an example. The supported algorithms and features for security and authentication are driven from the silicon cost and development cost. The algorithms chosen would also be those required by the IP storage standards. The authentication engine, block 2803, is illustrated to include the SHA-1

15 algorithm as one example of useable algorithms. This block provides message digest and authentication capabilities as specified in the IP security standards. The data flows through these blocks when security and message authentication services are required. The clear packets on their way out to the target are encrypted and are then authenticated if required using the appropriate engines. The secure packets received go through the same steps in

20 reverse order. The secure packet is authenticated and then decrypted using the engines 2803, 2804 of this block. The security engine also maintains the security associations in a security context memory, block 2809, that are established for the connections. The security associations (may include secure session index, security keys, algorithms used, current state of session and the like) are used to perform the message authentication and the

25 encryption/decryption services. It is possible to use the message authentication service and the encryption/decryption services independent of each other. The security engine of Fig. 28 or classification/policy engine of Fig. 20 or a combination thereof along with other protocol processing hardware capabilities of this patent create a secure TCP/IP stack using the hardware processor of this patent.

30 Fig. 29 illustrates the session cache and memory controller complex seen generally at 1704 of Fig. 17 in more detail. The memory complex includes a cache/memory architecture for the TCP/IP session database called session/global session memory or session cache in this patent, implemented as a cache or memory or a combination thereof. The session cache look-up engine, block 2904, provides the functionality to look-up a specific session cache

35 entry. This look-up block creates a hash index out of the fields provided or is able to accept a hash key and looks-up the session cache entry. If there is no tag match in the cache array

with the hash index, the look-up block uses this key to find the session entry from the external memory and replaces the current session cache entry with that session entry. It provides the session entry fields to the requesting packet processor complex. The cache entries that are present in the local processor complex cache are marked shared in the
5 global cache. Thus when any processor requests this cache entry, it is transferred to the global cache and the requesting processor and marked as such in the global cache. The session memory controller is also responsible to move the evicted local session cache entries into the global cache inside this block. Thus only the latest session state is available at any time to any requesters for the session entry. If the session cache is full, a new entry
10 may cause the least recently used entry to be evicted to the external memory. The session memory may be single way or multi-way cache or a hash indexed memory or a combination thereof, depending on the silicon real estate available in a given process technology. The use of a cache for storing the session database entry is unique, in that in networking applications for network switches or routers, generally there is not much locality of reference
15 properties available between packets, and hence use of cache may not provide much performance improvement due to cache misses. However, the storage transactions are longer duration transactions between the two end systems and may exchange large amounts of data. In this scenario or cases where a large amount of data transfer occurs between two nodes, like in clustering or media servers or the like a cache based session
20 memory architecture will achieve significant performance benefit from reducing the enormous data transfers from the off chip memories. The size of the session cache is a function of the available silicon die area and can have an impact on performance based on the trade-off. The memory controller block also provides services to other blocks that need to store packets, packet fragments or any other operating data in memory. The memory
25 interface provides single or multiple external memory controllers, block 2901, depending on the expected data bandwidth that needs to be supported. This can be a double data rate controller or controller for DRAM or SRAM or RDRAM or other dynamic or static RAM or combination thereof. The figure illustrates multi-controllers however the number is variable depending on the necessary bandwidth and the costs. The memory complex may also
30 provide timer functionality for use in retransmission time out for sessions that queue themselves on the retransmission queues maintained by the session database memory block.

Fig. 30 illustrates the data structures details for the classification engine. This is one way of organizing the data structures for the classification engine. The classification database is
35 illustrated as a tree structure, block 3001, with nodes, block 3003, in the tree and the actions, block 3008, associated with those nodes allow the classification engine to walk

down the tree making comparisons for the specific node values. The node values and the fields they represent are programmable. The action field is extracted when a field matches a specific node value. The action item defines the next step, which may include extracting and comparing a new field, performing other operations like ALU operations on specific data fields associated with this node-value pair, or may indicate a terminal node, at which point the classification of the specific packet is complete. This data structure is used by the classification engine to classify the packets that it receives from the packet scheduler. The action items that are retrieved with the value matches, while iterating different fields of the packet, are used by the results compiler to create a classification tag, which is attached to the packet, generally before the packet headers. The classification tag is then used as a reference by the rest of the processor to decide on the actions that need to be taken based on the classification results. The classifier with its programmable characteristics allows the classification tree structure to be changed in-system and allow the processor to be used in systems that have different classification needs. The classification engine also allows creation of storage /network policies that can be programmed as part of the classification tree-node-value-action structures and provide a very powerful capability in the IP based storage systems. The policies would enhance the management of the systems that use this processor and allow enforcement capabilities when certain policies or rules are met or violated. The classification engine allows expansion of the classification database through external components, when that is required by the specific system constraints. The number of trees and nodes are decided based on the silicon area and performance tradeoffs. The data structure elements are maintained in various blocks of the classification engine and are used by the classification sequencer to direct the packet classification through the structures. The classification data structures may require more or less fields than those indicated depending on the target solution. Thus the core functionality of classification may be achieved with fewer components and structures without departing from the basic architecture. The classification process walks through the trees and the nodes as programmed. A specific node action may cause a new tree to be used for the remaining fields for classification. Thus, the classification process starts at the tree root and progress through the nodes until it reaches the leaf node.

Fig. 31 illustrates a read operation between an initiator and target. The initiator sends a READ command request, block 3101, to the target to start the transaction. This is an application layer request which is mapped to specific SCSI protocol command which is then transported as an READ protocol data unit, block 3102, in an IP based storage network. The target prepares the data that is requested, block 3103 and provides read response PDUs, block 3105, segmented to meet the maximum transfer unit limits. The initiator then

retrieves the data, block 3016, from the IP packets and is then stored in the read buffers allocated for this operation. Once all the data has been transferred the target responds with command completion and sense status, block 3107. The initiator then retires the command once the full transfer is complete, block 3109. If there were any errors at the target and the
5 command is being aborted for any reason, then a recovery procedure may be initiated separately by the initiator. This transaction is a standard SCSI READ transaction with the data transport over IP based storage protocol like iSCSI as the PDUs of that protocol.

Fig. 32 illustrates the data flow inside the IP processor of this invention for one of the received READ PDUs of the transaction illustrated in Fig. 31. The internal data flow is
10 shown for the read data PDU received by the IP processor on the initiator end. This figure illustrates various stage of operation that a packet goes through. The stages can be considered as pipeline stages through which the packets traverse. The number of pipe stages traversed depends on the type of the packet received. The figure illustrates the pipe stages for a packet received on an established connection. The packet traverses through
15 the following major pipe stages:

1. Receive Pipe Stage of block 3201, with major steps illustrated in block 3207: Packet is received by the media access controller. The packet is detected, the preamble/ trailers removed and a packet extracted with the layer2 header and the payload. This is the stage where the Layer2 validation occurs for the intended recipient as well as any error
20 detection. There may be quality of service checks applied as per the policies established. Once the packet validation is clear the packet is queued to the input queue.

2. Security Pipe Stage of block 3202, with major steps illustrated in block 3208. The packet is moved from the input queue to the classification engine, where a quick determination for security processing is made and if the packet needs to go through security
25 processing, it enters the security pipe stage. If the packet is received in clear text and does not need authentication, then the security pipe stage is skipped. The security pipe stage may also be omitted if the security engine is not integrated with the IP processor. The packet goes through various stages of security engine where first the security association for this connection is retrieved from memory, and the packet is authenticated using the
30 message authentication algorithm selected. The packet is then decrypted using the security keys that have been established for the session. Once the packet is in clear text, it is queued back to the input queue controller.

3. Classification Pipe Stage of block 3203, with major steps illustrated in block 3209. The scheduler retrieves the clear packet from the input queue and schedules the

packet for classification. The classification engine performs various tasks like extracting the relevant fields from the packet for layer 3 and higher layer classification, identifies TCP/IP/ storage protocols and the like and creates those classification tags and may also take actions like rejecting the packet or tagging the packet for bypass depending on the policies
5 programmed in the classification engine. The classification engine may also tag the packet with the session or the flow to which it belongs along with marking the packet header and payload for ease of extraction. Some of the tasks listed may be or may not be performed and other tasks may be performed depending on the programming of the classification engine. As the classification is done, the classification tag is added to the packet and packet
10 is queued for the scheduler to process.

4. Schedule Pipe Stage of block 3204, with major steps illustrated in block 3210. The classified packet is retrieved from the classification engine queue and stored in the scheduler for it to be processed. The scheduler performs the hash of the source and destination fields from the packet header to identify the flow to which the packet belongs, if
15 not done by the classifier. Once the flow identification is done the packet is assigned to an execution resource queue based on the flow dependency. As the resource becomes available to accept a new packet, the next packet in the queue is assigned for execution to that resource.

5. Execution Pipe Stage of block 3205, with major steps illustrated in block
20 3211. The packet enters the execution pipe stage when the resource to execute this packet becomes available. The packet is transferred to the packet processor complex that is supposed to execute the packet. The processor looks at the classification tag attached to the packet to decide the processing steps required for the packet. If this is an IP based storage packet, then the session database entry for this session is retrieved. The database
25 access may not be required if the local session cache already holds the session entry. If the packet assignment was done based on the flow, then the session entry may not need to be retrieved from the global session memory. The packet processor then starts the TCP engine/ the storage engines to perform their operations. The TCP engine performs various TCP checks including checksum, sequence number checks, framing checks with necessary
30 CRC operations, and TCP state update. Then the storage PDU is extracted and assigned to the storage engine for execution. The storage engine interprets the command in the PDU and in this particular case identifies it to be a read response for an active session. It then verifies the payload integrity and the sequence integrity and then updates the storage flow state in the session database entry. The memory descriptor of the destination buffer is also
35 retrieved from the session data base entry and the extracted PDU payload is queued to the

storage flow/RDMA controller and the host interface block for them to DMA the data to the final buffer destination. The data may be delivered to the flow controller with the memory descriptor and the command/operation to perform. In this case deposit the data for this active read command. The storage flow controller updates its active command database.

- 5 The execution engine indicates to the scheduler the packet has been retired and the packet processor complex is ready to receive its next command.

6. DMA Pipe Stage of block 3206, with major steps illustrated in block 3212.

- Once the storage flow controller makes the appropriate verification of the Memory descriptor, the command and the flow state, it passes the data block to the host DMA engine for transfer
10 to the host memory. The DMA engine may perform priority based queuing, if such QOS mechanism is programmed or implemented. The data is transferred to the host memory location through DMA. If this is the last operation of the command, then the command execution completion is indicated to the host driver. If this is the last operation for a command and the command has been queued to the completion queue, the resources
15 allocated for the command are released to accept new command. The command statistics may be collected and transferred with the completion status as may be required for performance analysis, policy management or other network management or statistical purposes.

- Fig. 33 illustrates write command operation between an initiator and a target. The Initiator
20 sends a WRITE command, block 3301, to the target to start the transaction. This command is transported as a WRITE PDU, block 3302, on the IP storage network. The receiver queues the received command in the new request queue. Once the old commands in operation are completed, block 3304, the receiver allocates the resources to accept the WRITE data corresponding to the command, block 3305. At this stage the receiver issues a
25 ready to transfer (R2T) PDU, block 3306, to the initiator, with indication of the amount of data it is willing to receive and from which locations. The initiator interprets the fields of the R2T requests and sends the data packets, block 3307, to the receiver as per the received R2T. This sequence of exchange between the initiator and target continues until the command is terminated. A successful command completion or an error condition is communicated to the
30 initiator by the target as a response PDU, which then terminates the command. The initiator may be required to start a recovery process in case of an error. This is not shown in the exchange of the Fig. 33.

- Fig. 34 illustrates the data flow inside the IP processor of this invention for one of the R2T PDUs and the following write data of the write transaction illustrated in Fig. 33. The initiator
35 receives the R2T packet through its network media interface. The packet passes through all

the stages, blocks 3401, 3402, 3403, and 3404 with detailed major steps in corresponding blocks 3415, 3416, 3409 and 3410, similar to the READ PDU in Fig. 32 including Receive, Security, Classification, Schedule, and Execution. Security processing is not illustrated in this figure. Following these stages the R2T triggers the write data fetch using the DMA stage shown in Fig. 34, blocks 3405 and 3411. The write data is then segmented and put in TCP/IP packets through the execution stage, blocks 3406 and 3412. The TCP and storage session DB entries are updated for the WRITE command with the data transferred in response to the R2T. The packet is then queued to the output queue controller. Depending on the security agreement for the connection, the packet may enter the security pipe stage, block 3407 and 3413. Once the packet has been encrypted and message authentication codes generated, the packet is queued to the network media interface for the transmission to the destination. During this stage, block 3408 and 3414 the packet is encapsulated in the Layer 2 headers, if not already done so by the packet processor and is transmitted. The steps followed in each stage of the pipeline are similar to that of the READ PDU pipe stages above, with additional stages for the write data packet stage, which is illustrated in this figure. The specific operations performed in each stage depend on the type of the command, the state of the session, the command state and various other configurations for policies that may be setup.

Fig. 35 illustrates the READ data transfer using RDMA mechanism between an initiator and target. The initiator and target register the RDMA buffers before initiating the RDMA data transfer, blocks 3501, 3502, and 3503. The initiator issues a READ command, block 3510, with the RDMA buffer as the expected recipient. This command is transported to the target, block 3511. The target prepares the data to be read, block 3504, and then performs the RDMA write operations, block 3505 to directly deposit the read data into the RDMA buffers at the initiator without the host intervention. The operation completion is indicated using the command completion response.

Fig. 36 illustrates the internal architecture data flow for the RDMA Write packet implementing the READ command flow. The RDMA write packet also follows the same pipe stages as any other valid data packet that is received on the network interface. This packet goes through Layer 2 processing in the receive pipe stage, blocks 3601 and 3607, from where it is queued for scheduler to detect the need for security processing. If the packet needs to be decrypted or authenticated, it enters the security pipe stage, blocks 3602 and 3608. The decrypted packet is then scheduled to the classification engine for it to perform the classification tasks that have been programmed, blocks 3603 and 3609. Once classification is completed, the tagged packet enters the schedule pipe stage, blocks 3604 and 3610, where the scheduler

assigns this packet to a resource specific queue dependent on flow based scheduling. When the intended resource is ready to execute this packet, it is transferred to that packet processor complex, blocks 3605 and 3611, where all the TCP/IP verification, checks, and state updates are made and the PDU is extracted. Then the storage engine identifies the PDU as belonging to a storage flow for storage PDUs implemented using RDMA and interprets the RDMA command. In this case it is RDMA write to a specific RDMA buffer. This data is extracted and passed on to the storage flow/RDMA controller block which performs the RDMA region translation and protection checks and the packet is queued for DMA through the host interface, blocks 3606 and 3612. Once the packet has completed operation through the packet processor complex, the scheduler is informed and the packet is retired from the states carried in the scheduler. Once in the DMA stage, the RDMA data transfer is completed and if this is the last data transfer that completes the storage command execution, that command is retired and assigned to the command completion queue.

Fig. 37 illustrates the storage write command execution using RDMA Read operations. The initiator and target first register their RDMA buffers with their RDMA controllers and then also advertise the buffers to their peer. Then the initiator issues a write command, block 3701, to the target, where it is transported using the IP storage PDU. The recipient executes the write command, by first allocating the RDMA buffer to receive the write and then requesting an RDMA read to the initiator, blocks 3705, and 3706. The data to be written from the initiator is then provided as an RDMA read response packet, blocks 3707 and 3708. The receiver deposits the packet directly to the RDMA buffer without any host interaction. If the read request was for data larger than the segment size, then multiple READ response PDUs would be sent by the initiator in response to the READ request. Once the data transfer is complete the completion status is transported to the initiator and the command completion is indicated to the host.

Fig. 38 illustrates the data flow of an RDMA Read request and the resulting write data transfer for one section of the flow transaction illustrated in Fig. 37. The data flow is very similar to the write data flow illustrated in Fig. 34. The RDMA read request packet flows through various processing pipe stages including: receive, classify, schedule, and execution, blocks 3801, 3802, 3803, 3804, 3815, 3816, 3809 and 3810. Once this request is executed, it generates the RDMA read response packet. The RDMA response is generated by first doing the DMA, blocks 3805 and 3811, of the requested data from the system memory, and then creating segments and packets through the execution stage, blocks 3806 and 3812. The appropriate session database entries are updated and the data packets go to the security stage, if necessary, blocks 3807 and 3813. The secure or clear packets are then

queued to the transmit stage, block 3808 and 3814, which performs the appropriate layer 2 updates and transmits the packet to the target.

Fig. 39 illustrates an initiator command flow for the storage commands initiated from the initiator in more details. As illustrated following are some of the major steps that a command follows:

1. Host driver queues the command in processor command queue in the storage flow/RDMA controller;
2. Host is informed if the command is successfully scheduled for operation and to reserve the resources;
3. The storage flow/RDMA controller schedules the command for operation to the packet scheduler, if the connection to the target is established. Otherwise the controller initiates the target session initiation and once session is established the command is scheduled to the packet scheduler;
4. The scheduler assigns the command to one of the SAN packet processors that is ready to accept this command;
5. The processor complex sends a request to the session controller for the session entry;
6. The session entry is provided to the packet processor complex;
7. The packet processor forms a packet to carry the command as a PDU and is scheduled to the output queue; and
8. The command PDU is given to the network media interface, which sends it to the target.

This is the high level flow primarily followed by most commands from the initiator to the target when the connection has been established between an initiator and a target.

Fig. 40 illustrates read packet data flow in more detail. Here the read command is initially send using a flow similar to that illustrated in Fig. 39 from the initiator to the target. The target sends the read response PDU to the initiator which follows the flow illustrated in Fig. 40. As illustrated the read data packet passes through following major steps:

1. Input packet is received from the network media interface block;

2. Packet scheduler retrieves the packet from the input queue;
3. Packet is scheduled for classification;
4. Classified packet returns from the classifier with a classification tag;
5. Based on the classification and flow based resource allocation, the packet is assigned to a packet processor complex which operates on the packet;
6. Packet processor complex looks-up session entry in the session cache (if not present locally);
7. Session cache entry is returned to the packet processor complex;
8. Packet processor complex performs the TCP/IP operations / IP storage operations and extracts the read data in the payload. The read data with appropriate destination tags like MDL(memory descriptor list) is provided to the host interface output controller; and
9. The host DMA engine transfers the read data to the system buffer memory.

Some of these steps are provided in more details in Fig. 32, where a secure packet flow is represented, where as the Fig. 40 represents a clear text read packet flow. This flow and other flows illustrated in this patent are applicable to storage and non-storage data transfers by using appropriate resources of the disclosed processor, that a person with ordinary skill in the art will be able to do with the teachings of this patent.

Fig. 41 illustrates the write data flow in more details. The write command follows the flow similar to that in Fig. 39. The initiator sends the write command to the target. The target responds to the initiator with a ready to transfer (R2T) PDU which indicates to the initiator that the target is ready to receive the specified amount of data. The initiator then sends the requested data to the target. Fig. 41 illustrates the R2T followed by the requested write data packet from the initiator to the target. The major steps followed in this flow are as follows:

1. Input packet is received from the network media interface block;
2. Packet scheduler retrieves the packet from the input queue;
3. Packet is scheduled for classification;
4. Classified packet returns from the classifier with a classification tag;

a. Depending on the classification and flow based resource allocation, the packet is assigned to a packet processor complex which operates on the packet;

5. Packet processor complex looks-up session entry in the session cache (if not present locally);

5 6. Session cache entry is returned to the packet processor complex;

7. The packet processor determines the R2T PDU and requests the write data with a request to the storage flow/RDMA Controller;

8. The flow controller starts the DMA to the host interface;

9. Host interface performs the DMA and returns the data to the host input
10 queue;

10. The packet processor complex receives the data from the host input queue;

11. The packet processor complex forms a valid PDU and packet around the data, updates the appropriate session entry and transfers the packet to the output queue; and

15 12. The packet is transferred to the output network media interface block which transmits the data packet to the destination.

The flow in Fig. 41 illustrates clear text data transfer. If the data transfer needs to be secure, the flow is similar to that illustrated in Fig. 43, where the output data packet is routed through the secure packet as illustrated by arrows labeled 11a and 11b. The input R2T packet, if
20 secure would also be routed through the security engine (this is not illustrated in the figure).

Fig. 42 illustrates the read packet flow when the packet is in cipher text or is secure. This flow is illustrated in more details in Fig. 32 with its associated description earlier. The primary difference between the secure read flow and the clear read flow is that the packet is initially classified as secure packet by the classifier, and hence is routed to the security
25 engine. These steps are illustrated by arrows labeled 2a, 2b, and 2c. The security engine decrypts the packet and performs the message authentication, and transfers the clear packet to the input queue for further processing as illustrated by arrow labeled 2d. The clear packet is then retrieved by the scheduler and provided to the classification engine as illustrated by arrows labeled 2e and 3 in Fig. 42. The rest of the steps and operations are
30 the same as that in Fig. 40, described above.

Fig. 44 illustrates the RDMA buffer advertisement flow. This flow is illustrated to be very similar to any other storage command flow as illustrated in the Fig. 39. The detailed actions taken in the major steps are different depending on the command. For RDMA buffer advertisement and registration, the RDMA region id is created and recorded along with the address translation mechanism for this region is recorded. The RDMA registration also includes the protection key for the access control and may include other fields necessary for RDMA transfer. The steps to create the packet for the command are similar to those of Fig. 39.

Fig. 45 illustrates the RDMA write flow in more details. The RDMA writes appear like normal read PDUs to the initiator receiving the RDMA write. The RDMA write packet follows the same major flow steps as a read PDU illustrated in Fig. 40. The RDMA transfer involves the RDMA address translation and region access control key checks, and updating the RDMA database entry, beside the other session entries. The major flow steps are the same as the regular Read response PDU.

Fig. 46 illustrates the RDMA Read data flow in more details. This diagram illustrates the RDMA read request being received by the initiator from the target and the RDMA Read data being written out from the initiator to the target. This flow is very similar to the R2T response followed by the storage write command. In this flow the storage write command is accomplished using RDMA Read. The major steps that the packet follows are primarily the same as the R2T/write data flow illustrated in Fig. 41.

Fig. 47 illustrates the major steps of session creation flow. This figure illustrates the use of the control plane processor for this slow path operation required at the session initiation between an initiator and a target. This functionality is possible to implement through the packet processor complex. However, it is illustrated here as being implemented using the control plane processor. Both approaches are acceptable. Following are the major steps during session creation:

1. The command is scheduled by the host driver;
2. The host driver is informed that the command is scheduled and any control information required by the host is passed;
3. The storage flow/RDMA controller detects a request to send the command to a target for which a session is not existing, and hence it passes the request to the control plane processor to establish the transport session;

4. Control plane processor sends a TCP SYN packet to the output queue;
5. The SYN packet is transmitted to the network media interface from which is transmitted to the destination;
6. The destination, after receiving the SYN packet, responds with the SYN-ACK response, which packet is queued in the input queue on receipt from the network media interface;
7. The packet is retrieved by the packet scheduler;
8. The packet is passed to the classification engine;
9. The tagged classified packet is returned to the scheduler;
10. The scheduler, based on the classification, forwards this packet to control plane processor;
11. The processor then responds with an ACK packet to the output queue;
12. The packet is then transmitted to the end destination thus finishing the session establishment handshake; and
13. Once the session is established, this state is provided to the storage flow controller. The session entry is thus created which is then passed to the session memory controller (this part not illustrated in the figure).

Prior to getting the session in the established state as in step 13, the control plane processor may be required to perform a full login phase of the storage protocol, exchanging parameters and recording them for the specific connection if this is a storage data transfer connection. Once the login is authenticated and parameter exchange complete, does the session enter the session establishment state shown in step 13 above.

Fig. 48 illustrates major steps in the session tear down flow. The steps in this flow are very similar to those in Fig. 47. Primary difference between the two flows is that, instead of the SYN, SYN-ACK and ACK packets for session creation, FIN, FIN-ACK and ACK packets are transferred between the initiator and the target. The major steps are otherwise very similar. Another major difference here is that the appropriate session entry is not created but removed from the session cache and the session memory. The operating statistics of the connection are recorded and may be provided to the host driver, although this is not illustrated in the figure.

Fig. 49 illustrates the session creation and session teardown steps from a target perspective. Following are the steps followed for the session creation:

1. The SYN request from the initiator is received on the network media interface;
2. The scheduler retrieves the SYN packet from the input queue;
- 5 3. The scheduler sends this packet for classification to the classification engine;
4. The classification engine returns the classified packet with appropriate tags;
5. The scheduler, based on the classification as a SYN packet, transfers this packet to the control plane processor;
6. Control plane processor responds with a SYN-ACK acknowledgement packet.
- 10 It also requests the host to allocate appropriate buffer space for unsolicited data transfers from the initiator (this part is not illustrated);
7. The SYN-ACK packet is sent to the initiator;
8. The initiator then acknowledges the SYN-ACK packet with an ACK packet, completing the three-way handshake. This packet is received at the network media interface
- 15 and queued to the input queue after layer 2 processing;
9. The scheduler retrieves this packet;
10. The packet is sent to the classifier;
11. Classified packet is returned to the scheduler and is scheduled to be provided to the control processor to complete the three way handshake;
- 20 12. The controller gets the ACK packet;
13. The control plane processor now has the connection in an established state and it passes the to the storage flow controller which creates the entry in the session cache; and
14. The host driver is informed of the completed session creation.
- 25 The session establishment may also involve the login phase, which is not illustrated in the Fig. 49. However, the login phase and the parameter exchange occur before the session enters the fully configured and established state. These data transfers and handshake may

primarily be done by the control processor. Once these steps are taken the remaining steps in the flow above may be executed.

Figs. 50 and 51 illustrate write data flow in a target subsystem. The Fig. 50 illustrates an R2T command flow, which is used by the target to inform the initiator that it is ready to accept a data write from the initiator. The initiator then sends the write which is received at the target and the internal data flow is illustrated in Fig. 51. The two figures together illustrate one R2T and data write pairs. Following are the major steps that are followed as illustrated in Figs. 50 and 51 together:

1. The target host system in response to receiving a write request like that illustrated in Fig. 33, prepares the appropriate buffers to accept the write data and informs the storage flow controller when it is ready, to send the ready to transfer request to the initiator;
2. The flow controller acknowledges the receipt of the request and the buffer pointers for DMA to the host driver;
3. The flow controller then schedules the R2T command to be executed to the scheduler;
4. The scheduler issues the command to one of the packet processor complexes that is ready to execute this command;
5. The packet processor requests the session entry from the session cache controller;
6. The session entry is returned to the packet processor;
7. The packet processor forms a TCP packet and encapsulates the R2T command and sends it to the output queue;
8. The packet is then sent out to network media interface which then sends the packet to the initiator. The security engine could be involved, if the transfer needed to be secure transfer;
9. Then as illustrated in Fig. 51, the initiator responds to R2T by sending the write data to the target. The network media interface receives the packet and queues it to the input queue;
10. The packet scheduler retrieves the packet from the input queue;

11. The packet is scheduled to the classification engine;
 12. The classification engine provides the classified packet to the scheduler with the classification tag. The flow illustrated is for unencrypted packet and hence the security engine is not exercised;
 - 5 13. The scheduler assigns the packet based on the flow based resource assignment queue to packet processor queue. The packet is then transferred to the packet processor complex when the packet processor is ready to execute this packet;
 14. The packet processor requests the session cache entry (if it does not already have it in its local cache);
 - 10 15. The session entry is returned to the requesting packet processor;
 16. The packet processor performs all the TCP/IP functions, updates the session entry and the storage engine extracts the PDU as the write command in response to the previous R2T. It updates the storage session entry and routes the packet to the host output queue for it to be transferred to the host buffer. The packet may be tagged with the memory
15 descriptor or the memory descriptor list that may be used to perform the DMA of this packet into the host allocated destination buffer; and
 17. The host interface block performs the DMA, to complete this segment of the Write data command.
- Fig. 52 illustrates the target read data flow. This flow is very similar to the initiator R2T and
20 write data flow illustrated in Fig. 41. The major steps followed in this flow are as follows:
1. Input packet is received from the network media interface block;
 2. Packet scheduler retrieves the packet from the input queue;
 3. Packet is scheduled for classification;
 4. Classified packet returns from the classifier with a classification tag;
 - 25 a. Depending on the classification and flow based resource allocation, the packet is assigned to a packet processor complex which operates on the packet
 5. Packet processor complex looks-up session entry in the session cache (if not present locally);

6. Session cache entry is returned to the packet processor complex;
7. The packet processor determines the Read Command PDU and requests the read data with a request to the flow controller;
8. The flow controller starts the DMA to the host interface;
- 5 9. Host interface performs the DMA and returns the data to the host input queue;
- 10 10. The packet processor complex receives the data from the host input queue;
11. The packet processor complex forms a valid PDU and packet around the data, updates the appropriate session entry and transfers the packet to the output queue;
- 10 and
12. The packet is transferred to the output network media interface block which transmits the data packet to the destination.

The discussion above of the flows is an illustration of some the major flows involved in high bandwidth data transfers. There are several flows like fragmented data flow, error flows with
15 multiple different types of errors, name resolution service flow, address resolution flows, login and logout flows, and the like are not illustrated, but are supported by the IP processor of this invention.

As discussed in the description above, the perimeter security model is not sufficient to protect an enterprise network from security threats due to the blurring boundary of enterprise
20 networks. Further, a significant number of unauthorized information access occurs from inside. The perimeter security methods do not prevent such security attacks. Thus it is critical to have security deployed across the network and protect the network from within as well as the perimeter. The network line rates inside enterprise networks are going to 1Gbps, multi-Gbps and 10Gbps in the LANs and SANs. As previously mentioned, distributed firewall
25 and security methods require a significant processing overhead on each of the system host CPU if implemented in software. This overhead can cause increase in latency of the response of the servers, reduce their overall throughput and leave fewer processing cycles for applications. An efficient hardware implementation that can enable deployment of software driven security services is required to address the issues outlined above. The
30 processor of this patent addresses some of these key issues. Further, at high line rates it is critical to offload the software based TCP/IP protocol processing from the host CPU to protocol processing hardware to reduce impact on the host CPU. Thus, the protocol

processing hardware should provide the means to perform the security functions like firewall, encryption, decryption, VPN and the like. The processor provides such a hardware architecture that can address the growing need of distributed security and high network line rates within enterprise networks.

- 5 Fig. 53 illustrates a traditional enterprise network with perimeter firewall. This figure illustrates local area network and storage area networks inside enterprise networks. The figure illustrates a set of clients, 5301(1) through 5301(n), connected to an enterprise network using wireless LAN. There may be multiple clients of different types like handheld computers, PCs, thin clients, laptops, notebook computers, tablet PCs and the like. Further,
- 10 they may connect to the enterprise LAN using wireless LAN access points (WAP), 5303. There may be one or more WAP connected to the LAN. Similarly, the figure also illustrates multiple clients connected to the enterprise LAN through wired network. These clients may be on different sub segments or the same segment or be directly linked to the switches in a point to point connection, depending on the size of the network, the line rates and the like.
- 15 The network may have multiple switches and routers that provide the internal connectivity for the network of devices. The figure also illustrates network attached storage devices, 5311, providing network file serving and storage services to the clients. The figure also illustrates one or more servers, 5307(1) through 5307(n) and 5308(1) through 5308(n), attached to the network providing various application services being hosted on these servers to the clients
- 20 inside the network as well as those being accessed through the outside as web access or other network access. The servers in the server farm may be connected in a traditional three-tier or n-tier network providing different services like web server, application servers, database servers, and the like. These servers may hold direct attached storage devices for the needed storage and/or connect to a storage area network (SAN), using SAN connectivity
- 25 and switches, 5309(1) through 5309(n) to connect to the storage systems, 5310(1) through 5310(n) for their storage needs. The storage area network may also be attached to the LAN using gateway devices, 5313 to provide the access to storage system to the LAN clients. The storage systems may also be connected to the LAN directly, similar to NAS, 5311, to provide block storage services using protocols like iSCSI and the like. This is not illustrated
- 30 in the figure. The network illustrated in this figure is secured from the external network by the perimeter firewall, 5306. As illustrated in this figure the internal network in such an environment does not enable security, which poses serious security vulnerabilities to insider attacks.

- Figure 54 illustrates an enterprise network with a distributed firewall and security capabilities.
- 35 The network configuration illustrated is similar to that in Fig. 53. The distributed security

features shown in such a network may be configured, monitored, managed, enabled and updated from a set of central network management systems by central IT manager(s), 5412. The manager(s) is(are) able to set the distributed security policy from management station(s), distribute appropriate policy rules to each node enabled to implement the distributed security policy and monitor any violations or reports from the distributed security processors using the processor of this patent. The network may be a network that comprises of one or more nodes, one or more management stations or a combination thereof. The figure illustrates that the SAN devices are not under the distributed security network. The SAN devices in this figure may be under a separate security domain or may be trusted to be protected from insiders and outsiders with the security at the edge of the SAN.

Figure 55 illustrates an enterprise network with a distributed firewall and security capabilities where the SAN devices are also under a distributed security domain. The rest of the network configuration may be similar to that in Fig. 54. In this scenario, the SAN devices may implement similar security policies as the rest of the network devices and may be under the control from the same IT management systems. The SAN security may be implemented different from the rest of the network, depending on the security needs, sensitivity of the information and potential security risks. For instance, the SAN devices may implement full encryption/decryption services beside firewall security capabilities to ensure that no unauthorized access occurs as well as the data put out on the SAN is always in a confidential mode. These policies and rules may be distributed from the same network management systems or there may be special SAN management systems, not shown, that may be used to create such distributed secure SANs. The systems in Fig. 54 and fig 55 use the processor and the distributed security system of this patent.

Fig. 56 illustrates a central manager/policy server and monitoring station, also called the central manager. The central manager includes security policy developer interface, block 5609, which is used by the IT manager(s) to enter the security policies of the organization. The security policy developer interface may be a command line interface, a scripting tool, a graphical interface or a combination thereof which may enable the IT manager to enter the security policies in a security policy description language. It may also provide access to the IT manager remotely under a secure communication connection. The security policy developer interface works with a set of rule modules that enables the IT manager to enter the organization's policies efficiently. The rule modules may provide rule templates that may be filled in by the IT managers or may be interactive tools that ease the entry of the rules. These modules provide the rules based on the capabilities that are supported by the

distributed security system. Networking layers 2 through 4 (L2, L3, L4) rules, rule types, templates, and the like is provided by block 5601 to the security developer interface. These rules may comprise of IP addresses for source, destination, L2 addresses for source, destination, L2 payload type, buffer overrun conditions, type of service, priority of the connection, link usage statistics and the like or a combination thereof. The Protocol/port level rules, block 5602, provides rules, rule types, templates and the like to the security developer interface. These rules may comprise of protocol type like IP, TCP, UDP, ICMP, IPSEC, ARP, RARP or the like, or source port, or destination port including well-known ports for known upper level applications/protocols, or a combination thereof. The block 5603 provides application level or upper layer (L5 through L7) rules, rule types, templates and the like to the security developer interface. These rules may comprise rules that are dependent on a type of upper layer application or protocol like HTTP, XML, NFS, CIFS, iSCSI, iFCP, FCIP, SSL, RDMA or the like, their usage model, their vulnerabilities or a combination thereof. The content based rules, block 5604, provide rules, rule types, templates, or the like to the security developer interface for entering content dependent rules. These rules may evolve over time, like the other rules, to cover known threats or potential new threats and comprise of a wide variety of conditions like social security numbers, confidential/proprietary documents, employee records, patient records, credit card numbers, offending URLs, known virus signatures, buffer overrun conditions, long web addresses, offending language, obscenities, spam, or the like or a combination thereof. These rules, templates or the rule types may be provided for ease of creation of rules in the chosen policy description language(s) for the manager of the distributed security system. Security policy developer interface may exist without the rules modules and continue to provide means to the IT managers to enter the security policies in the system. The rules represented in the security policy language entered through the interface would then get compiled by the security rules compiler, block 5611, for distribution to the network nodes. Security rules compiler utilizes a network connectivity database, 5605, and a nodes capabilities and characteristics database, 5606, to generate rules specific for each node in the network that is part of monitoring/enforcing the security policy. The network connectivity database comprises physical adjacency information, or physical layer connectivity, or link layer connectivity, or network layer connectivity, or OSI layer two addresses or OSI layer three addresses or routing information or a combination thereof. The nodes capabilities and characteristics database comprises hardware security features or software security features or size of the rules engine or performance of the security engine(s) or quality of service features or host operating system or hosted application(s) or line rates of the network connectivity or host performance or a combination thereof. The information from these databases would enable the security rules compiler to properly map security policies to node specific rules. The node

specific rules and general global rules are stored to and retrieved from the rules database, 5607. The security rules compiler then works with the rules distribution engine, 5608, to distribute the compiled rules to each node. The rules distribution engine interacts with each security node of the distributed security system to send the rule set to be used at that
5 specific node. The rule distribution engine may retrieve the rule sets directly from the rules database or work with the security rules compiler or a combination thereof to retrieve the rules. Once the rules are proliferated to respective nodes the central manager starts monitoring and managing the network.

The central manager works with each node in the security network to collect events or
10 reports of enforcement, statistics, violations and the like using the event and report collection/management engine, 5616. The event/report collection engine works with the security monitoring engine, 5613, to create the event and information report databases, 5614 and 5615, which keep a persistent record of the collected information. The security monitoring engine analyzes the reports and events to check for any violations and may in
15 turn inform the IT managers about the same. Depending on the actions to be taken when violations occur, the security monitoring engine may create policy or rule updates that may be redistributed to the nodes. The security monitoring engine works with the security policy manager interface, 5612, and policy update engine, 5610, for getting the updates created and redistributed. The security policy manager interface provides tools to the IT manager to
20 do event and information record searches. The IT manager may be able to develop new rules or security policy updates based on the monitored events or other searches or changes in the organizations policies and create the updates to the policies. These updates get compiled by the security policy compiler and redistributed to the network. The functionality of security policy manager interface, 5612, and policy update engine, 5610, may be provided
25 by the security policy developer interface, 5609, based on an implementation choice. Such regrouping of functionality and functional blocks is possible without diverging from the teachings of this patent. The security monitoring engine, the security policy manager interface and the event/report collection/management interface may also be used to manage specific nodes when there are violations that need to be addressed or any other actions
30 need to be taken like enabling a node for security, disabling a node, changing the role of a node, changing the configuration of a node, starting /stopping/deploying applications on a node, or provisioning additional capacity or other management functions or a combination thereof as appropriate for the central manager to effectively manage the network of the nodes.

Fig. 57 illustrates the central manager flow of this patent. The central manager may comprise various process steps illustrated by the blocks of the flow. The IT manager(s) create and enter the security policies of the organization in central management system(s) that are illustrated by block 5701. The policies are then compiled into rules, by the security policy compiler, using a network connectivity database and a node capabilities and characteristics database as illustrated by block 5702. The central manager then identifies the nodes from the network that have security capability enabled, from the node characteristics database, in block 5703, to distribute rules to these nodes. The manager may then select a node from these nodes, as illustrated by block 5704, and retrieve the corresponding security rules from the rules database, as illustrated by block 5705, and then communicate the rules to the node, as illustrated by 5706, and further illustrated by Fig. 58. The central manager continues the process of retrieving the rules and communicating the rules until all nodes have been processed as illustrated by the comparison of all nodes done in block 5707. Once rules have been distributed to all the nodes, the central manager goes into managing and monitoring the network for policy enforcements, violations or other management tasks as illustrated by block 5708. If there are any policy updates that result from the monitoring, the central manager exits the monitoring to create and update new policy through checks illustrated by blocks, 5709 and 5710. If there are new policy updates, the central manager traverses through the flow of Fig. 57 to compile the rules and redistribute them to the affected nodes and then continue to monitor the network. The event collection engine of the central manager continues to monitor and log events and information reports, when other modules are processing the updates to the security policies and rules. Thus the network is continuously monitored when the rule updates and distribution is in progress. Once the rule updates are done, the security monitoring engine and other engines process the collected reports. Communication of rules to the nodes and monitoring/managing of the nodes may be done in parallel to improve the performance as well as effectiveness of the security system. Central manager may communicate new rules or updates to multiple nodes in parallel instead of using a serial flow, and assign the nodes that have already received the rules into monitoring/managing state for the central manager. Similarly the policy creation or updates can also be performed in parallel to the rule compilation, distribution and monitoring.

Fig. 58 illustrates the rule distribution flow of this patent. The rule distribution engine working with the security policy compiler, retrieves the rules or rule set to be communicated to a specific node as illustrated by 5801. It then initiates communication with the selected node as illustrated by 5802. The central manager and the node may authenticate each other using agreed upon method or protocol as illustrated by 5803. Authentication may involve a

complete login process, or secure encrypted session or a clear mode session or a combination thereof. Once the node and the central managers authenticate each other, the communication is established between the central manager and the control plane processor or host based policy driver of the node as illustrated by 5804. Once the communication is established, the rule distribution engine sends the rules or rule set or updated rules or a combination thereof to the node as illustrated in 5805. This exchange of the rules may be over a secure/encrypted session or clear link dependent on the policy of the organization. The protocol deployed to communicate the rules may be using a well known protocol or a proprietary protocol. Once the rule set has been sent to the node, the central manager may wait to receive the acknowledgement from the node of successful insertion of the new rules at the node as illustrated by 5806. Once a successful acknowledgement is received the rule distribution flow for one node concludes as illustrated by 5807. The appropriate rule database entries for the node would be marked with the distribution completion status. The flow of Fig. 58 is repeated for all nodes that need to receive the rules from the rule distribution engine of the central manager. The rule distribution engine may also be able to distribute rules in parallel to multiple nodes to improve the efficiency of the rule distribution process. In this scenario the rule distribution engine may perform various steps of the flow like authenticate a node, establish communication with a node, send rule or rules to a node and the like in parallel for multiple nodes.

Fig. 59 illustrates a control plane processor or a host based policy driver flow of this patent. This flow is executed on each node following the distributed security of this patent, comprising a hardware processor. Upon initiation of policy rule distribution by the central manager or upon reset or power up or other management event or a combination thereof the policy driver establishes communication with the central manager/policy server as illustrated by 5901. The policy driver receives the rule set or updates to existing rules from the central manager as illustrated by 5902. If the rules are formatted to be inserted into the specific policy engine implementation, size and the like, the rules are accepted to be configured in the policy engine. If the rules are always properly formatted by the central manager it is feasible to avoid performing the check illustrated in block 5903. Otherwise, if the rules are not always formatted or otherwise ready to be directly inserted in the policy engine, as determined in block 5903, the driver configures the rules for the node as illustrated by block 5904. The driver then communicates with the database initialization and management interface, block 2011 of Fig. 20, of the policy engine of the processor. This is illustrated by block 5905. Then the driver sends a rule to the policy engine which updates it in the engine data structures, like that in Fig. 30, which comprises of a ternary or binary CAM, associated memory, ALU, database description and other elements in the classification/policy engine of

Fig. 20. This is illustrated by block 5906. This process continues until all the rules have been entered in the policy engine through the decision process illustrated by 5907, 5908 and 5906. Once all rules have been entered, the policy engine activates the new rules working with the driver as illustrated by block 5909. The driver then updates/sends the rules to a
5 persistent storage for future reference and/or retrieval as illustrated by block 5910. The driver then communicates to the central manager/policy server of the update completion and new rules activation in the node as illustrated by block 5911. The policy driver may then enter a mode of communicating the management information, events, reports to the central manager. This part of the driver is not illustrated in the figure. The management
10 functionality may be taken up by a secure process on the host or the control plane processor of the node. The mechanisms described above allow a secure operating environment to be created for the protocol stack processing, where even if the host system gets compromised either through a virus or malicious attack, it allows the network security and integrity to be maintained since a control plane processor based policy driver does not allow the host
15 system to influence the policies or the rules. The rules that are active in the policy engine would prevent a virus or intruder to use this system or node to be used for further virus proliferation or attacking other systems in the network. The rules may also prevent the attacker from extracting any valuable information from the system like credit card numbers, social security numbers, medical records or the like. This mechanism significantly adds to
20 the trusted computing environment needs of the next generation computing systems.

Some or all portions of the flow may be performed in parallel as well as some portions may be combined together. For instance, one or more rules may be communicated together by the policy driver to the database initialization/management interface, block 2011, which may then update the rules in the policy engine in an atomic fashion instead of doing it one rule at
25 a time. Further, while new rules are being received by the policy driver or the policy engine or a combination thereof, the hardware processor may continue to perform rule enforcement and analysis with the active rule set in parallel on the incoming or outgoing network traffic.

Fig. 60 illustrates rules that may be deployed in a distributed security system using this patent. The IT manager(s) may decide the policies that need to be deployed for different
30 types of accesses. These policies are converted into rules at the central management system, 5512 or 5412, for distribution to each node in the network that implements one or more security capabilities. The rules are then provided to the processor on the related node. A control plane processor, 1711 of Fig. 17, working with classification and policy engine, 1703, and the DB Initialization/management control interface, 2011 of Fig. 20, of the
35 processor configure the rule in the processor. Each node implementing the distributed

security system may have unique rules that need to be applied on the network traffic passing through, originating or terminating at the node. The central management system interacts with all the appropriate nodes and provides each node with its relevant rules. The central management system also interacts with the control plane processor which works with the classification/policy engine of the node to retrieve rule enforcement information and other management information from the node for distributed security system.

Fig. 60 illustrates rules that may be applicable to one or more nodes in the network. The rules may contain more or fewer fields than indicated in the figure. In this illustration, the rules comprise the direction of the network traffic to which the rule is applicable, either In or Out; the source and destination addresses, which may belong to an internal network node address or address belonging to a node external to the network; protocol type of the packet, e.g TCP, UDP, ICMP and the like as well as source port and destination ports and any other deep packet fields comprising URL information, sensitive information like credit card numbers or social security numbers, or any other protected information like user names, passwords and the like. The rule then contains an action field that indicates the action that needs to be taken when a certain rule is matched. The action may comprise of various types like permit the access, deny the access, drop the packet, close the connection, log the request, send an alert or combination of these or more actions as may be appropriate to the rule matched. The rules may be applied in a priority fashion from top to bottom or any other order as may be implemented in the system. The last rule indicates a condition when none of the other rules match and, as illustrated in this example, access is denied.

Fig 61 illustrates TCP/IP processor version of the IP processor illustrated in Fig. 16 and Fig. 17. This processor consists of a network interface block 6101, which is used to connect this processor to the network. The network interface may be a wired or wireless Ethernet interface, Packet over Sonet interface, Media Independent Interface (MII), GMII, XGMII, XAUI, System Packet Interface, SPI 4 or SPI 5 or other SPI derivatives or other network protocol interface or a combination thereof. This is the interface used to send or receive packets to or from the network to which this processor is coupled. Intelligent flow controller and packet buffer block 6103, provides packet scheduler functionality of block 1702 of Fig. 17 as well as the input and output queue controller functionality of block 1701 and 1712. Programmable classification/Rule Engine/Security Processing block 6102, provides the classification and policy/rule processing functionality of block 1703 as well as the security processing functionality of the block 1705 when security capabilities are supported by the specific implementation.

TCP/IP packet processor engines of block 6104, are similar to the TCP/IP processor engine of SAN packet processor of blocks 1706(a) through 1706(n). The connection (session) memory block 6105, provides the functionality of IP session cache/memory of block 1704, whereas the connection manager and control plane processor of block 6106 provide the session controller and control plane processor functionality similar to that of blocks 1704 and 1711. The RDMA controller block 6107, provides RDMA functionality similar to the block 1708. The memory controller block 6109, provides memory interface similar to that provided by memory controller of block 1704. The TCP/IP processor may have external memory which may be SRAM, DRAM, FLASH, ROM, EEPROM, DDR SDRAM, RDRAM, FCRAM, QDR SRAM, Magnetic RAM or Magnetic memory or other derivatives of static or dynamic random access memory or a combination thereof. Host/Fabric/Network Interface block 6108 provides the interface to a host bus or a switch fabric interface or a network interface depending on the system in which this processor is being incorporated. For example in a server or server adapter environment the block 6108 would provide a host bus interface functionality similar to that of block 1710, where the host bus may be a PCI bus, PCI-X, PCI-Express, or other PCI derivatives or other host buses like AMBA bus, or RapidIO bus or HyperTransport or other derivatives. A switch or a router or a gateway or an appliance with a switch fabric to connect multiple line cards would have appropriate fabric interface functionality for block 6108. This may include queues with priority mechanisms to avoid head of the line blocking, fragmentation and defragmentation circuitry as needed by the switch fabric, and appropriate flow control mechanism to ensure equitable usage of the switch fabric resources. In case of an environment like a gateway or appliance that connects to a network on ingress and egress, the block 6108 would provide network interface functionality similar to the block 6101.

The TCP/IP processor illustrated in Fig. 61 is a version of the architecture shown in Fig. 16 and Fig. 17 as is evident from the description above. The TCP/IP processor engines of block 6104 may be substituted with SAN packet processors of block 1706(a) through 1706(n) and the two architectures would offer the same functionality. Thus Fig. 61 can be looked at as a different view and/or grouping of the architecture illustrated in Fig. 16 and Fig. 17. The TCP/IP processor engines may be augmented by the packet engine block of the SAN packet processors to provide programmable processing where additional services can be deployed besides protocol processing on a packet by packet basis. Block 6110 of Fig. 61, illustrated as a dotted line around a group of blocks, is called "TCP/IP processor core" in this patent. The RDMA block 6107 is shown to be part of the TCP/IP Processor core although it is an optional block in certain TCP/IP processor core embodiments like low line speed applications or applications that do not support RDMA. Similarly the security engine

may also not be present depending on the implementation chosen and the system embodiment.

Fig. 62 illustrates an Adaptable TCP/IP processor of this patent. This processor comprises of a network interface block 6201, host/fabric/network interface block 6207, a TCP/IP processor core block 6202, a runtime adaptable processor (RAP) block 6206, or a combination thereof. The adaptable TCP/IP processor may also include an adaptation controller block 6203, configuration memory block 6204, a memory interface block 6205, data buffers block 6209, a memory controller block 6208, RAP interface block 6210, RAP Extension interface block 6211, or a combination thereof. The TCP/IP processor core, block 6202, is the TCP/IP processor core illustrated in Fig. 61 block 6110. As discussed earlier the security and RDMA blocks of the TCP/IP processor core may or may not be present depending on the application and system environment. The TCP/IP processor core provides full TCP/IP protocol processing, protocol termination and protocol initiation functionality. The TCP/IP processor core may provide TCP/IP protocol stack comprising at least one of the following hardware implemented functions:

- a. sending and receiving data, including upper layer data;
- b. establishing transport sessions and session teardown functions;
- c. executing error handling functions;
- d. executing time-outs;
- e. executing retransmissions;
- f. executing segmenting and sequencing operations;
- g. maintaining protocol information regarding active transport sessions;
- h. maintaining TCP/IP state information for each of one or more session connections.
- i. fragmenting and defragmenting data packets;
- j. routing and forwarding data and control information;
- k. sending to and receiving from a peer, memory regions reserved for RDMA;

- l. recording said memory regions reserved for RDMA in an RDMA database and maintaining said database;
- m. executing operations provided by RDMA capability;
- n. executing security management functions;
- 5 o. executing policy management and enforcement functions;
- p. executing virtualization functions;
- q. communicating errors;
- r. processing Layer 2 media access functions to receive and transmit data packets, validate the packets, handle errors, communicate errors and other Layer 2 functions;
- 10 s. processing physical layer interface functions;
- t. executing TCP/IP checksum generation and verification functions;
- u. processing Out of Order packets;
- v. CRC calculation functions;
- 15 w. processing Direct Data Placement/Transfer;
- x. Upper Layer Framing functions;
- y. processing functions and interface to socket API's;
- z. forming packet headers for TCP/IP for transmitted data and extraction of payload from received packets; and
- 20 aa. processing header formation and payload extraction for Layer 2 protocols of data to be transmitted and received data packets; respectively.

The TCP/IP processor core may provide a transport layer RDMA capability as described earlier. The TCP/IP processor core may also provide security functions like network layer security, transport layer security, socket layer security, application layer security or a combination thereof besides wire speed encryption and decryption capabilities. Thus the TCP/IP processor core may also provide a secure TCP/IP stack in hardware with several

functions described above implemented in hardware. Even though the description of the adaptable TCP/IP processor has been with the TCP/IP processor core as illustrated in this application, the TCP/IP processor core may have various other architectures. Beside the architecture disclosed in this patent, the TCP/IP processor core could also be a fixed
5 function implementation, or may be implemented as a hardware state machine or may support partial protocol offloading capability for example support fast path processing in hardware where control plane processing as well as session management and control may reside in a separate control plane processor or host processor or a combination of various architecture alternatives described above. The TCP/IP processor core architecture chosen
10 may also include functions for security or RDMA or a combination thereof. Further, the adaptable TCP/IP processor architecture can be used for protocols other than TCP/IP like SCTP, UDP or other transport layer protocols by substituting the TCP/IP processor core with a protocol appropriate processor core. This would enable creating an adaptable protocol processor targeted to the specific protocol of interest. The runtime adaptable processor of
15 such a processor would be able to function similarly to the description in this patent and offer hardware acceleration for similar applications/services by using its dynamic adaptation capabilities.

The runtime adaptable processor, block 6206, provides a dynamically changeable hardware where logic and interconnect resources can be adapted programmatically on the fly to create
20 virtual hardware implementations as appropriate to the need of the application/service. The adaptation controller, block 6203, may be used to dynamically update the RAP block. The adaptation controller may interface with the host processor or control plane processor or the TCP/IP processor or a combination thereof to decide when to switch the configuration of RAP block to create a new avatar or incarnation to support needed hardware function(s),
25 what function(s) should this avatar of RAP block perform, where to fetch the new avatar, how long is the avatar valid, when to change the avatar, as well as provide multiple simultaneous function support in the RAP block. The RAP block may be dynamically switched from one avatar to another avatar, depending on the analysis done in TCP/IP processor core. For instance, the TCP/IP processor core may have a programmed policy that will ask it to flag
30 any data payload received that may contain XML data and pass the extracted data payload for processing through the RAP instead of sending it directly to the host processor. In this instance, when a packet is received that contains XML data, the TCP/IP processor core may tag the data appropriately and either queue the packets in the external memory for further processing by RAP or pass the data in the data buffers of block 6209 for further processing
35 by RAP. The TCP/IP processor core may be coupled to a RAP interface, block 6210, which may provide the functionality needed for the TCP/IP processor core to interface with the

RAP and the adaptation controller block. This functionality may be directly part of RAP or adaptation controller or the TCP/IP processor core. The RAP interface would inform the adaptation controller in this instance of the arrival of XML traffic, so the adaptation controller can fetch the appropriate configuration from configuration memory, block 6204, which may
5 be internal or external memory or a combination thereof. The adaptation controller can then provide the configuration to RAP block 6206 and enable it when the XML data is ready to be operated on and is ready in the data buffers or external memory for the RAP to fetch it. Similarly, depending on the policies that may be programmed in the TCP/IP processor core and the received network traffic, the RAP block may get configured into a new hardware
10 avatar to support the specific application, service or function or a combination thereof dynamically based on the characteristics of the received traffic and the policies.

The TCP/IP processor core may also choose to pass the received data to the host processor without passing it for further processing through the RAP depending on the policies and/or the nature of the data received. Thus, if hardware configurations for specific operations or
15 functions or policies or applications have not been realized, because the operations or policies or functions or applications may not be used often and hence do not cause performance issues or resources have not been assigned to develop the acceleration support for cost reasons or any other business or other reasons, those operations may be performed on the host. As those operations are realized as a runtime adaptable
20 configuration, it may be provided to the adaptation controller so it can configure the RAP block for that operation as needed dynamically. The TCP/IP processor would also be informed to then identify such operations and pass them through RAP block. Using such a technique, over a period of time more applications can be accelerated without the need for changing or adding any hardware accelerators. The deployment of new policies or
25 applications or services or operations on the runtime adaptable processor may be under the user or administrator control using very similar mechanisms as those shown for the security policy deployment and management. Thus, a central administrator can efficiently deploy new configurations to systems with the runtime adaptable protocol processor of this patent as and when needed. Similarly, the user or the administrator could remove or change the
30 RAP supported functions or policies or applications as the need or the usage of the system changes. For example, a system using runtime adaptable protocol processor of this patent may initially be used for XML traffic, however its usage may change to support voice over IP application and XML acceleration may not be required, but instead some other voice over IP acceleration is needed. In such an instance the user or the administrator may be able to
35 change, add or remove selectable hardware supported configurations from the specific system or systems. The central manager/policy server flow, central manager flow, rule

distribution flow, control plane processor/policy driver flows and the like illustrated in Fig. 56 through Fig. 59 are applicable to the management, deployment, change, monitoring and the like for the runtime adaptable configurations as well with appropriate changes similar to those explained as follows. The runtime adaptable configuration creation flow may be added
5 to the security policy creation flow for example. New configurations may become available from another vendor and the user may just need to select the configuration of interest to be deployed. The configuration distribution flow may be similar to the rule distribution flow, where the policies for the support of the configuration(s) may be distributed to the TCP/IP processor core blocks, where as the configuration may be distributed to the adaptation
10 controller of the system of interest or to a driver or an configuration control process on the host system or a combination thereof. Thus the runtime adaptable protocol processor systems may be integrated well into other enterprise management systems when used in that environment. The application or policy or service or operation configurations may be distributed by other means for example as a software update over the network or through
15 mass storage devices or other means. The foregoing description is one way of providing the updates in one usage environment but there can multiple other ways to do the same for each embodiment and the usage environment as one skilled in the art can appreciate and hence should not be viewed as limited to the description above.

The adaptation controller may also be required to configure the RAP block to operate on
20 data being sent out to the network. In such a case the RAP block may be required to operate on the data before it is passed on to TCP/IP processor core to send it to the intended recipient over the network. For example, it may be necessary to perform secure socket layer (SSL) operations on the data before being encapsulated in the transport and network headers by the TCP/IP processor core. The host driver or the application that is
25 sending this data would inform the adaptation controller of the operation to be performed on the data before being passed on to the TCP/IP processor core. This can happen through the direct path from the host/fabric interface 6207 to the adaptation controller 6203. The adaptation controller can then configure RAP block 6206 or a part of it to perform the operation requested dynamically and let RAP operate on the data. Once RAP operation is
30 completed it can inform the adaptation controller of the operation completion, which can then work with the TCP/IP processor core to send this data enroute to its destination after appropriate protocol processing, header encapsulation and the like by the TCP/IP protocol processor. RAP 6206 may pass the processed data to the TCP/IP processor core through data buffers of block 6209 or by queuing them in memory using the memory interface block
35 6205. Thus the runtime adaptable TCP/IP processor of this patent can be configured to

operate on incoming as well as outgoing data, before or after processing by the TCP/IP processor core.

Runtime adaptable processor 6206 may be restricted in size or capabilities by physical, cost, performance, power or other constraints. RAP extension interface, block 6211, may also be provided on the adaptable TCP/IP processor to interface RAP block 6206 to one or more external components providing runtime adaptable processor functionality. Thus the solution can be scaled to bigger size or features or capabilities using the RAP extension interface 6211. RAP extension interface comprises of all the necessary control, routing, data, memory interface buses and connections as needed to seamlessly extend the RAP into one or more external components.

Fig. 63 illustrates an adaptable TCP/IP processor alternative of this patent to that described above. As indicated earlier, the TCP/IP processor portion of this processor may not only be the architecture disclosed in this patent but may also be a fixed function implementation, or may be implemented as a hardware state machine or may support partial protocol offloading capability, for example support fast path processing in hardware where control plane processing as well as session management and control may reside in a separate control plane processor or host processor or a combination of various architecture alternatives described above. The TCP/IP processor core architecture chosen may also include functions for security or RDMA or a combination thereof. Further, the adaptable TCP/IP processor architecture can be used for protocols other than TCP/IP like SCTP, UDP or other transport layer protocols by substituting the TCP/IP processor core with a protocol appropriate processor core. The adaptable TCP/IP processor alternate of Fig. 63 illustrates the runtime adaptable processor, block 6311, the adaptation controller, block 6310, and the configuration memory, block 6312, as integrated more tightly in the TCP/IP processor architecture to create a runtime adaptable TCP/IP processor. The functions provided by RAP, block 6311, adaptation controller, block 6310, and the configuration memory, block 6312, is very similar to that of the corresponding blocks in Fig. 62. The RAP interface functionality of block 6210, or the memory interface block 6205, or data buffers, block 6209, may be appropriately provided by blocks 6310, 6311 or 6312 or a combination thereof. It may also be distributed within the TCP/IP processor elements. This architecture may also provide a RAP extension interface like that of block 6211 to provide RAP scalability, even though such a block is not shown in Fig. 63. This version of the adaptable TCP/IP processor would also operate similar to that in Fig. 62 and can also be configured to operate on incoming as well as outgoing data, before or after processing by the TCP/IP processor core blocks.

Fig. 64 illustrates a runtime adaptable processor of this patent. The runtime adaptable processor comprises computational logic and interconnect resource that can be dynamically changed to map various hardware functions that need to be accelerated. The computational logic blocks may be realized using FPGA like combinational blocks for fine grain control or

5 may consist of one or more simple programmable processor(s), ALU, and memory that can be configured to provide specific hardware function(s) at a given time which may then be changed dynamically to support a new function. The dynamic change of the function can be done by a configuration controller as needed by the usage of RAP. For example the computational block(s) may be setup to provide addition operation for a selected time period

10 on incoming data to the computational block, but then as a new avatar is created the operation provided may be selected to be subtraction for the duration of the new avatar. The selection of the new operation for the new avatar may be done by the appropriate configuration controller. Thus the function provided by a computational block can be dynamically changed to another function, while some other computational blocks may

15 continue to provide their selected operation. The computational block function change may take a certain period of time, which may be as low as a clock period or multiple clock periods or other period. The dynamic adaptation of one or more computational blocks may be done simultaneously or otherwise as needed. The interconnect resources may also be realized similarly to that of reconfigurable routing resources of FPGAs. Fig. 64 illustrates a runtime

20 adaptable processor architecture of this patent as a hierarchy of computational logic, called compute clusters, blocks 6401(1) through 6401(Z), interconnected using a routing network comprised of routing resources 6407(a) through 6407(n). These routing resources are interconnected using the inter cluster routing switch, 6403. The inter cluster routing switch may be configured dynamically to provide highly programmable interconnections between

25 various compute clusters thereby creating changing avatars of the hardware. The compute clusters may be configured individually by the global configuration controller, block 6405, which works with the adaptation controller, block 6203 of Fig 62, to dynamically adapt the RAP. The global configuration controller works with configuration memory controller, block 6406 and configuration memory block 6204 and the adaptation controller 6203, both of Fig.

30 62, to retrieve configuration information which is used to dynamically change the individual compute cluster configurations and the inter cluster routing switch for interconnect configurations. Input/Output interface and controller, block 6404, is used to interface the runtime adaptable processor with adaptation controller, block 6203, data buffers, block 6209, RAP extension interface, block 6211 or the host/fabric/network interface, block 6207, all of

35 Fig. 62. The global memory and controller, block 6402, provides global memory to compute clusters and also provides a controller to interface with external memory interface block 6205. Computational logic inside compute clusters 6401(1) – 6401(Z) may need memory

beside that inside each cluster. The global memory block and controller can fulfill this need. The figure illustrates multiple interconnection elements that serve different roles.

Interconnect channels 6407(a) through 6407(n), are the routing channels to connect each cluster to the inter cluster routing switch to enable a multi-way connection capability for each cluster to source or sink information from other clusters of the RAP. Interconnect channels 6408(a) through 6408(n) provide memory interconnect resources for the compute clusters for them to get access to the global memory. These memory channels may be shared among the clusters in a column or there may be multiple paths to the memory which may be used simultaneously by many compute clusters in a column to write or read data to or from the memory. Interconnect channels 6409(a) through 6409(m) are the configuration channels that allow the global configuration controller to send configuration information to the compute clusters and receive event information or other information from compute clusters that may be used to change the configuration of a given cluster or some of the clusters or all RAP cluster configurations. The interconnect channel architecture and implementation for the above may be accomplished using wide busses, high speed serial interconnects or other implementation choices. The specific choice or topology is not dictated or implied by the figure.

The runtime adaptable processor of Figure 64 can be configured such that computation array may be split into partial regions, where each region may be configured to perform a specific hardware operation. For example clusters 6401(1), 6401(2) may form one region whereas clusters 6401(3) through 6401(M) may form another region and some other clusters may form yet another region. Some of the regions may be interconnected as pipelined stages as may be required by the hardware function being mapped onto the runtime adaptable processor. Regions of the mapping may interconnect with each other or may operate on independent data or streams as may be appropriate for the operations mapped. The regions can all be dynamically adapted with the changing needs of the processing requirements. The regions can be very granular or may involve only partial compute clusters as well. Hence the runtime adaptable processor of this patent is dynamically adaptable to a very fine grain level to meet the demands of the required processing.

Fig. 65 illustrates a compute cluster of this patent. The compute cluster comprises computational elements (CE), blocks 6501(1) through 6501(Z), that provide computational logic. CEs may be composed of FPGA-like combinational logic blocks and interconnect resources or may be simple programmable processors with ALU and memory which provides a given hardware function based on the instructions programmed. CEs may be dynamically configured by changing the instruction being executed on the input data or

stored data or combination thereof to perform a new hardware function. The processors may be simple processors supporting hard wired instructions through combinational logic that can select required hardware operation configured in the combinational logic. The processors may be more complex processors where the hardware configuration may select the instruction that is executed through the ALU and other functional resources providing a virtual hardware avatar/incarnation/configuration. The avatar may also comprise multiple instructions being executed through the resources forming a more complex configuration. Multiple avatars may be programmed in the CE, and a specific avatar can be dynamically selected, providing very flexible hardware architecture. The CEs may provide bit-wise operations, as well as operations on groups of bits like 4-bit, 8-bit, 16-bit groupings as desired by the granularity of the configuration options. This bit groupings may be an implementation choice where larger or smaller groupings can be selected without deviating from the principles of the teachings of this patent. The cluster configuration controller, block 6507, interacts with the global configuration controller, block 6405 of Fig. 64, to select the specific avatar for each CE. The interconnect channels, 6502(1) through 6502(N), provide the configuration information from the configuration controller to the CEs and any execution information or events or avatar change requests or a combination thereof from the CEs to the cluster configuration controller. This information may be used to direct the flow of the configuration controller mapping different avatars in conjunction with the global configuration controller 6405 and/or the adaptation controller 6203 of Fig. 62. This information may be used to create pipeline stages of operation where different portions of compute cluster or compute clusters provide multiple stages of operations. Interconnect channels, 6503(1) through 6503(M) provide connectivity to cluster memory and controller, block 6506, to the CEs. There may be multiple parallel paths into the cluster memory, thereby allowing multiple simultaneous accesses to different regions of the memory as indicated. CEs on a given channel may all share the channel or there may be multiple paths per channel to the memory as well. The cluster memory may be a single memory array or may be multiple memory arrays as an implementation choice. The cluster memory is also coupled to the global memory and controller, block 6402 of Fig. 64, through channels like 6508 and 6408(a) through 6408(n) of Fig. 64. The global memory and cluster memory may each be accessible from the host/fabric/network interface, 6207, or the adaptation controller, 6203, or memory interface, 6205 or a combination thereof to read or write memory locations individually or as a set of locations for initialization or other purposes like DMA access, test, or the like. The CEs may also provide connectivity to their next neighbor as indicated in Fig. 65 by the arrows. Not all neighbor connections indicated have to be present. This can be an implementation choice. These connections allow CEs to send or receive output or input data, flags, exception conditions or like information or a combination thereof to their

neighbors. The avatar selected for the CE would decide which inputs to use to retrieve the needed information to operate on. The outputs from CEs may also be selected to be coupled to the cluster routing switch, 6506, which can then provide selected connectivity between CEs within the cluster, as well as provide connectivity with other compute clusters by coupling with the inter cluster routing switch, 6403, and the interconnect channels like 6407(a) through 6407(n). The cluster routing switch may be configured for the appropriate interconnections through the cluster configuration controller, 6507, by coupling with interconnect channel, 6504.

Fig. 66 illustrates a security solution using the teachings of this patent. The security solution comprises a central manager, a network, one or more line cards and secure chips. The central manager is a collection of functional modules that reside in a central management system used by IT manager(s) to create, deploy and monitor security rules. Central manager modules are similar to those of the central manager shown in Fig. 56 and both are used interchangeably in the following description. These modules may reside on the same set of systems that are used for managing the overall network or may be on independent systems. Block 6601 is illustrated to represent network management applications and security applications that may be deployed for a network. These applications are used by the IT managers to create their security rules or policies. The central manager provides an application programmer interface (API), block 6602, which provides a uniform interface to security and management applications of 6601, to use the distributed security system of this patent. The API interfaces with network layer rules engine, block 6603, application layer rules engine, block 6604, storage area network rules engine, block 6605 or other application specific rule engines, block 6619, or a combination thereof. These rule engines provide API support functions for one or more specific categories of the rules that they represent. They may also provide rule templates that are preconfigured that an IT manager can use by filling in relevant fields of the rules for their specific needs. For instance there may be a set of rules that deny connection requests to all users whose network address is not a local address, or deny requests to specific ports like port 80 for all outside connections or the like or a combination thereof. These rules engines assemble the rules and provide them to the rules compiler, block 6606, for compiling them for distribution to secure nodes. The compiler uses nodes capability and connectivity database, block 6617, to compile node appropriate rules and actions. The compiled rules are deposited in compiled rules database, block 6618. The rules distribution engine, block 6607, distributes the rules to the appropriate nodes using a central manager flow and rules distribution flow similar to that illustrated in Fig. 57 and Fig. 58. The security rules may be distributed to the host processor or a control plane processor as illustrated in Fig 58 or to a control processor and scheduler, block 7103, described below,

or a combination thereof as appropriate depending on the node capability. The rules may be distributed using a secure link or insecure link using proprietary or standard protocols as appropriate per the specific node's capability over a network. The network may be a local area network (LAN), wide area network (WAN), metro area network (MAN), wireless LAN, storage area network (SAN) or a system area network or another network type deployed or a combination thereof. The network may be Ethernet based, internet protocol based or SONET based or other protocol based or a combination thereof. Monitoring interface, block 6609, and Event recording engine and database, block 6608, are utilized to collect various security and/or management events from various nodes that are being monitored for security violations or other conditions as defined by the rules. These blocks represent the central manager blocks 5613, 5614, 5615 and 5616 described above and provide similar functionality. The monitoring engine may provide the analysis capability as described for block 5613 or may work with analysis and reporting application(s) illustrated by block 6610 to provide intelligent reports to the IT manager of security violations or breaches or conformance or other issues upon request or automatically as programmed by the IT manager depending on the nature of the issue and its severity. The central manager modules may also be deployed local to a network node system, for example a switch or a router, and work within the system's control and management software. It may be used to deploy and monitor rules local to various line card(s) or accelerator cards or other cards providing security capability of the system. In such an instance, the network used to communicate the rules may be a local bus or a system area network, or a combination thereof, of the specific system.

Security Solution comprises line cards which may incorporate the security processor, SAN protocol processor, TCP/IP processor or runtime adaptable protocol processor or various other processors disclosed in this patent. The line card configuration and the architecture may vary with the specific system and the application. Three types of line card architectures, a) flow-through b) look-aside and c) accelerator card, are illustrated in this patent to illustrate usage models for the processors of this patent. Fig. 68, Fig. 69 and Fig. 70 illustrate these configurations using a security processor based system, though it could also be based on other processors of this patent. Blocks 6612 and block 6613 illustrate two of these types of card configurations. The security processor illustrated in these cards is that disclosed in this patent. There are various different variations of the security processor that can be created depending on the functionality incorporated in the processor. Blocks 6614 and block 6615 illustrate two versions of such security processor. Block 6614 illustrates a security processor core comprising at least a content search and rule processing engine coupled with a runtime adaptable processor. This processor is similar to that illustrated in

Fig. 71 and is described in detail below. Block 6615, illustrates the security processor of block 6614 coupled with a TCP/IP processor or a protocol processor to provide more functionality usable in a security node as a security processor. A reduced functionality security processor, not illustrated, may also be created by removing runtime adaptable processor and associated logic from block 6614 to provide a content search and rules processing engine based security processor. The choice of the security processor may depend on the system in which it is being deployed, the functionality supported by the system, the solution cost, performance requirement, or other reasons, or a combination thereof. The security processor may use one or more ports to connect to external memories, block 6616, which may be used to store rules information, or other intermediate data or packets or other information as necessary to perform various functions needed for security processing. The memories may be of various types like DRAM, SDRAM, DDR DRAM, SRAM, RDRAM, FCRAM, QDR SRAM, DDR SRAM, Magnetic memories, Flash or a combination thereof or future derivatives of such memory technologies. The inventions disclosed in this patent enable many variations of the architectures illustrated and may be appreciated by those skilled in the art that changes in the embodiments may be made without departing from the principles and spirit of the invention.

Fig. 67 illustrates security solution compiler flow. As described above security rules may be of various types like application layer rules, block 6701, network layer rules, block 6702, Storage area network rules, block 6703, or application specific rules, block 6619, or a combination thereof. As illustrated in Fig. 67 application layer rules comprise basic string search rules that may be expressed in a special language or a standard representation like regular expressions or a combination thereof. Application layer rules which typically require searching character strings deep inside a packet may be represented using a regular expression. The types of application layer rules or network layer rules or SAN rules or application specific rules may vary with the specific node where they may be deployed, the organization or the entity using them, the security threats being defended against or other purposes or a combination thereof. The figure illustrates various categories of rules that may be created depending on the usage model. These rules may be created by anti-spam software vendors or the entity using the security system or vendor supplying the security solution or other third parties or a combination thereof. There may be application layer rules to defend against spam. This may comprise of rules that have been created using the knowledge of spam or unwanted messages. For instance a rule may be to search for a message like "receive million USD" inside any incoming email anywhere within the email including the header and the message. Such a rule may be represented using a regular expression like `".*receive million [U|u][s|S][D|d]"` which will detect the message of interest i.e

"receive million USD" but may also detect variations of this where USD is not all capitals e.g. usd or UsD or usD or the like. The leading "." in this rule indicates to search for the message anywhere within the received data packets. A set of rules may be defined like the one above to form the anti-spam rule set. These rules may be updated as new types of spam or methods are discovered and can be kept up to date with constantly evolving threats. Similarly a set of rules may be developed to perform virus scan functions to detect for various known viruses. Anti-virus rules are typically signature matching or pattern matching rules similar to those discussed above. The virus signatures may be looked for at specific locations in a message or a file and may be described using a similar method. A set of anti-virus rules are defined from known virus signatures to detect for known viruses. As new viruses or worms become known, the anti-virus rules may be updated to defend against them as well. These rules would then be compiled through the security compiler flow and distributed to all nodes of interest as discussed earlier. Once these rules get deployed, the security nodes may be programmed to take action corresponding to the match on an anti-spam or anti-virus rule to deny access to the particular node originating the message or, drop the connection or flag the session to the IT manager or other appropriate action as defined by the rule. The central manager modules provide the ability for the IT manager to define such actions when certain conditions like those above are met. The actions may be comprised of one or more of drop connection, deny access, report violation, page the network manager, allow access but record the violation for later analysis or isolate the source node to a specific virtual LAN or transfer the connection to some other node or other similar action as appropriate.

Similarly there may be many other categories of application layer rules. For example there may be rules defined to manage digital rights of the owners of the electronic documents or media or content as may be appropriate. Such rules may also be defined similar to the signature or pattern matching or string of character matching rules above. These rules may flag matches to a specific digital rights signature inside a content, which can then be used to refer to a digital rights database that may indicate if such an access or usage of the digital content is permitted to the owner. The digital rights ownership data base may reside in the memories associated with the security processor and a control processor, like block 6809 or block 7103, described below, can refer to that database to decide if valid ownership exists or not and if it does not exist what specific action should be taken based on the defined rule. The digital rights confirmation may be done by some other device or processor in the specific node which is performing the digital rights signature matching. The decision of where to perform such analysis functionality may depend on the specific system usage

model and the system design choices. A set of rules for digital rights management also be created as part of the application layer rules for the security processor.

Instant messaging (IM) has gained tremendous success in its usage by individuals as well as corporations. Instant messaging may be regulated for various industries like the financial industry to preserve for future reference and is also subject to spam like other modes of communication like email. Thus some organizations may create rules specifically targeted towards instant messaging to protect against ensuing liabilities in case of wrongful usage or protect the users from unwanted spam or for other reasons as deemed appropriate by the organization. One of the issues with instant messaging is that any level of policing has to be done in stream without creating delays in the communication. Thus a hardware based security enforcement of this patent may be needed to monitor IM. These rules are similar to other application layer rules discussed above and may be created using similar means like defining the message search strings using regular expressions.

Recent surveys by FBI and others have found that over 70% of attacks on information technology are from within an organization. Thus there is a need for a class of security devices and rules that need to be developed to protect from the damaging effects of such attacks. These rules are defined as extrusion detection rules. The extrusion detection rules may be created to detect intentional or unintentional disclosure of confidential or proprietary or sensitive information of the organization using the network from going outside the perimeter of the organization. For example a software company may need to guard its core software source code from accidental or malicious disclosure to people or entities unauthorized to get it. A set of rules may thus be created by the organization that may search for specific strings or paragraphs or code modules or other appropriate information within all outbound messages and flag them or prevent them from being sent. Such rules may also be compiled using the security compiler flow and distributed to the appropriate node or nodes. For example a rule may be defined to search for a "Top Secret" phrase in any message being sent that is outbound from the organization and flag such a message for further review by the IT manager or to drop such connection and inform the user or other responsible person. A regular expression rule ".*Top Secret" may be defined to search for the term anywhere in a message. Such rules may also be created as application layer rules that may then be compiled and distributed to appropriate nodes for detection and enforcement of extrusion detection security functionality.

The IT manager may be able to create classes of rules from the application layer rules or network layer rules or SAN rules or application specific rules or other rules and deploy a class of rules to a class of security nodes and a different class of rules to another set of

security nodes. For example the manager can create certain application layer rules like anti-spam or anti-virus rules and network layer rules that are deployed to the switches and routers of the network that are security enabled with the teaching of this patent and another set of rules like intrusion detection rules and network layer rules for sensitive servers

5 holding critical top secret information. It may be possible to create different sets of rules that may be deployed depending on the functions within an organization. For example, security nodes that are deployed within a manufacturing department may get one set of rules while those in an engineering department may get a different set of rules. Creating a different set of rules for different types of devices or different device users or node specific rules or a

10 combination thereof can be used as a process to create a pervasive and layered security within an organization.

Similarly there may be application layer rules that detect or flag access to specific web address or URL's or other confidential information like customer information comprising their credit card numbers, or health information or financial reports or the like, which may be used

15 to create a different set of application rules as shown in block 6701. With an increase in usage of voice over IP solutions within organizations and over the internet, security threats are also increasing. It may then be necessary to create rules specific to VOIP, for example rogue connections may need to be detected and flagged or VOIP traffic may not be allowed to go outside an organization's boundary or detect for viruses entering the organization

20 through VOIP connections or create confidentiality of VOIP traffic by encrypting it or the like. The VOIP rules may also be created using the same application layer rules engines and detect matches to the rules at appropriate nodes in the network. The runtime adaptable processor, block 7102, described below, may be used to provide encryption or decryption services to VOIP traffic when such traffic is detected by the VOIP rule match. Similarly,

25 other application specific rules may also be developed and provided in the central manager modules to be programmed, compiled and distributed to the secure nodes in the network using the compiler flow illustrated in Fig. 67.

Network layer rules, block 6702, may comprise various rules targeted at the network and transport layers of the network. These rules are similar to those illustrated in Fig. 60. These

30 rules may include IP level address rules, protocol port rules, protocol specific rules, connection direction oriented rules, and the like. These rules may be described in a special language or using regular expressions. In TCP/IP based networks these are primarily TCP and IP header fields based rules, where matches may be defined on source address or destination address or an address range or port numbers or protocol type or a combination

35 thereof. Similarly there may be rules targeted specifically to storage area networks which

may transport critical information assets of an organization. This is shown as a different category of rules, but may comprise storage network's network layer rules, application layer rules or the like. There may be rules targeted to specific logical unit numbers (LUNs) or zones (groups of source/destination addresses) or logical or physical block addresses or the like. These rules may also be represented in a specific language or as strings of characters or data patterns using regular expressions.

The secure solution compiler of Fig. 67 allows an IT manager to create security rules of different types as discussed above and enable them to create a layered and/or pervasive security model. The compiler flow would be provided with the characteristics of the specific nodes like the security capability presence, the rules communication method, the size of the rule base supported, the performance metrics of the node, deployment location e.g. LAN or SAN or other, or the like. The compiler flow then uses this knowledge to compile node specific rules from the rule set(s) created by the IT manager. The compiler comprises a rules parser, block 6704, for parsing the rules to be presented to the lexical analyzer generator, block 6705, which analyzes the rules and creates rules database used for analyzing the content. The rule parser may read the rules from files of rules or directly from the command line or a combination depending on the output of the rule engines. The rules for a specific node are parsed to recognize the language specific tokens used to describe the rules or regular expression tokens. The parser then presents the tokens to the lexical analyzer generator. The lexical analyzer processes the incoming tokens and generates non-deterministic finite automaton (NFA) which represents rules for parsing the content. The NFA is then converted in deterministic finite automaton (DFA) by the lexical analyzer generator to enable deterministic processing of the rule states. The process of creating NFAs and DFAs is well understood by modern compiler developers. However, the lexical analyzer generator creates various tables that represent DFA states and the state transition tables for the rules that are used by a hardware lexical analyzer instead of generating lexical analysis software as is done for compilers. One way to view the rules is that they define a language to recognize the content. These tables are used by a lexical analyzer hardware or content search and rule processing engine, block 7106, described below, to analyze the stream of data being presented to the security processor of this patent. The regular expression rules can be viewed as defining a state transition table. For example, if a string "help" is being searched, using a regular expression "help", then each character of the regular expression can be viewed to represent a state. There may be a start state s0, and character specific states s1(h), s2(e), s3(l), and s4(p) where s(x) represent a state for a character x. There may also be error states like s_err which may be entered upon terminating a search when appropriate transition conditions are not met. As the input stream

is being analyzed by the hardware lexical analyzer this state machine is activated when a first "h" is encountered, and the state machine reaches s1. Now if the next character in the stream is an "e" then the state machine transitions to s2. Thus if a string "help" is encountered the state machine will reach state s4. States s1 through s3 are accepting states, meaning they continue the search to the next state. State s4, for this string is marked by the lexical analyzer generator as a terminal state. These states are marked as accepting or terminal states in the accept tables. When a comparison reaches a terminal state, a match with the specific rule may be indicated. Any action that needs to be taken based on matching of a rule is created in a match/action table as an action tag or instruction that is then used by the content search and rule processing engine, block 7106, to take specific action or forward the match and action information to control processor, block 7103, to take appropriate rule specific action. However, if there is only a partial rule match e.g. if the input content includes string "her", then the rule processing hardware will enter state s2, having encountered "he" however, as soon as "r" is analyzed, an error is indicated to mean that there is no rule match and processing of the input stream starts from that point forward from the initial state s0. Though the above description is given with regards to using single character match per state, it is possible to analyze multiple characters at the same time to speed up the hardware analysis. For example, the lexical analyzer generator may create tables that enable transition of 4 characters per state there by quadrupling the content search speed. The lexical analyzer generator creates character class tables, block 6706, next state look-up tables, block 6709, state transition tables, block 6707, accept states, block 6708 and match/action tables, block 6710 which are then stored in the compiled rules database storage, block 6711. The character class tables are created by compressing the characters that create a similar set of state transition into a group of states for compact representation. The state transition tables comprise of rows of states in a DFA table with compressed character class as the columns to look-up the next state transitions. The next state table are used to index to the next state from the current state in the state machine represented by the DFA. These tables are stored in on-chip and off-chip memories associated with security processors of this patent. The compiler of this patent uses the node characteristics and connectivity database to create the rules on a node by node basis. The compiler indicates an error to the IT manager if certain rules or rule sizes do not match the capabilities of the specific nodes so they may be corrected by the manager. This information is retrieved from a node characteristics and connectivity database as illustrated by block 6713.

Rules distribution engine, block 6712, follows the central manager and rules distribution flow illustrated in Fig. 57 and Fig. 58. The security rules may be distributed to the host processor

or a control plane processor as illustrated in Fig 58 or to a control processor and scheduler, block 7103, described below, or a combination thereof as appropriate depending on the node capability. The rules may be distributed using a secure link or insecure link using proprietary or standard protocols as appropriate per the specific node's capability over a network.

Fig. 71 illustrates a security processor of this patent. The security processor comprises a coprocessor or host bus interface, block 7101, a control processor and scheduler, block 7103, at least one content search and rules processing engine, block 7106, next state memory, block 7110, match/action table memory, block 7111, character class table memory, block 7107, and accept and state transition memories, block 7108. The security processor may also comprise of packet buffers, block 7104, memory controller, block 7112, run time adaptable processor, block 7102, adaptation controller, block 7105 and configuration memory, block 7109. A version of security processor may be created by using coprocessor or host interface controller acting as a data interface, a control processor and scheduler, at least one content search and rules processing engine, next state memory, match/action table memory, character class table memory, accept and state transition memories and memory controller. Memory controller may not be required in system applications where the number of rules is small enough to fit in the on chip memories. Such a processor may perform all the content search tasks; however it may not be able to provide targeted application acceleration, which may be feasible with a security processor that includes a run time adaptable processor.

The control processor and scheduler, block 7103, communicates with the rules distribution engine, block 6712 to receive appropriate data tables prior to starting the content inspection. It stores the received state information into their respective dedicated memories. The character class table from block 6706, is stored in the memory block 7107. The state transition and accept tables, block 6707 and 6708, are stored in their respective memories represented by block 7108. Block 7108 may also be two or more separate memories for performance reasons but are illustrated by one block in the figures. The next state look-up tables from block 6709 are stored in the next state memory, block 7110. The match/action tables from block 6710 are stored in their memory block 7111. These tables may be larger than the memory available in the security processor on-chip, and may be stored in external memory or memories that are accessed by the memory controller block 7112. There may be multiple ports to memory to speed up access to data tables stored in external memories. These memories may be of various types like DRAM, SDRAM, DDR DRAM, SRAM, RDRAM, FCRAM, QDR SRAM, DDR SRAM, Magnetic memories, Flash or a combination

thereof or future derivatives of such memory technologies. For most applications next state table and action tables may need to be off-chip, whereas the other tables may be maintained on chip dependent on the size and number of the rules. Once the rules distribution engine provides the tables to the control processor and scheduler, block 7103, and they are setup in
5 their respective memories, the security processor is ready to start processing the data stream to perform content inspection and identify potential security rule matches or violations. The security processor state configuration information is received via a coprocessor/host interface controller. The security processor of this patent may be deployed in various configurations like a look-aside configuration illustrated in Fig. 69 or flow-through
10 configuration illustrated in Fig. 68 or an accelerator adapter configuration illustrated in Fig. 70 as well others not illustrated which can be appreciated by persons skilled in the art. In a look-aside or an accelerator adapter configuration, the security processor of this patent is under control of a master processor which may be a network processor or a switch processor or a TCP/IP processor or classification processor or forwarding processor or a
15 host processor or the like depending on the system in which such a card would reside. The control processor and scheduler receives the configuration information under the control of such master processor that communicates with the rule engine to receive packets that contain the configuration information and passes it on to the security processor. Once the configuration is done, the master processor provides packets to the security processor for
20 which content inspection needs to be performed using the coprocessor or host interface. The coprocessor or the host interface may be standard buses like PCI, PCI-X, PCI express, RapidIO, HyperTransport or LA-1 or SRAM memory interface or the like or a proprietary bus. The bandwidth on the bus should be sufficient to keep the content search engine operating at its peak line rate. The security processor may be a memory mapped or an IO mapped
25 device in the master processor space for it to receive the packets and other configuration information in a look-aside or accelerator configuration. The security processor may be polled by the master processor or may provide a doorbell or interrupt mechanism to the master to indicate when it is done with a given packet or when it finds a match to the programmed rules. The control processor and scheduler, block 7103 and the block 7101
30 work with the master processor to provide the above functionality. The control processor and scheduler stores incoming packets to the packet buffer, block 7104, and schedules the packets for processing by the content search and rule processing engines as they become available to analyze the content. The scheduler maintains the record of the packets being processed by the specific engines and once the packets are processed it informs the master
35 processor. The content search and rule processing engines of block 7106 inform the control processor and the scheduler when they have found a match to a rule and the action associated with that rule as programmed in the match/action table. This information may in

turn be sent by the control processor to the master processor, where the master processor can take specific action for the packet indicated by the rule. The actions may be one from a multitude of actions like dropping the packet or dropping a connection or informing the IT manager, or the like, as discussed earlier. When the security processor includes a runtime

5 adaptable processor like block 7102, the control processor and scheduler may schedule operations on the packet through block 7102. The control processor would work with the adaptation controller, block 7105, to select the specific avatar of the processor for the needed operation. For example, a packet that needs to be decrypted before being analyzed

10 may be scheduled to the adaptable processor before being analyzed by the content search engines. Once the packet has been decrypted by the adaptable processor it is then scheduled by block 7103 to block 7106. However, runtime the adaptable processor may operate on a packet once a match has been found by the content search engines or the packet has been processed by the search engine without any issues. For example, the packet data may need to be encrypted once no issues have been found. The control

15 processor and scheduler schedules the packets to the runtime adaptable processor in the appropriate order as defined by the needs of the operation. The runtime adaptable processor, block 7102, adaptation controller, block 7105 and configuration memory, block 7109 is similar to those illustrated in figures 62, 63, 64 and 65. The runtime adaptable processor and the associated block provide similar functionality with appropriate logic

20 enhancements made to couple to the control processor and scheduler of the security processor. The runtime adaptable processor may be used to provide compression and decompression service to the packets if the appropriate adaptation configurations are deployed. The runtime adaptable processor may also be used for VOIP packets providing relevant hardware acceleration service to those packets like DSP processing or encryption

25 or decryption or the like.

The security processor may also need to provide inspection ability across multiple packets in a connection between a source and a destination. The control processor and scheduler, block 7103, provides such functionality as well. The control processor may store the internal processing state of the content search and security processing engine in a connection

30 database which may be maintained in the on chip memory in the control processor or in the off-chip memory. The control processor and scheduler looks up the execution or analysis state for a given connection when a packet corresponding to the connection is presented to it by the master processor or in the incoming traffic in a flow-through configuration described below. The connection ID may be created by the master processor and provided that to the

35 security processor with the packet to be inspected or the security device may derive the connection association from the header of the packet. The connection ID may be created in

the IP protocol case by using a 5-tuple hashing derived from the source address, destination address, source port, destination port and the protocol type. Once the connection ID is created and resolved in case of a hash conflict by the control processor and scheduler, it then retrieves the state associated with that connection and provides the state to the search engines, block 7106, to start searching from that state. This mechanism is used to create multi-packet searches per connection and detect any security violations or threats that span packet boundaries. For example, if there is a rule defined to search for "Million US Dollars" and if this string appears in a connection data transfer in two separate packets where "Million U" appears in one packet and "S Dollars" appears in another packet then if a connection based multi-packet search mechanism of this patent is not present the security violation may not be detected since each packet individually does not match the rule. However, when the multi-packet search is performed, no matter how far apart in time these two packets arrive at the security node, the state of the search will be maintained from one packet to another for the connection and the strings of two packets will be detected and flagged as a continuous string "Million US Dollars".

As discussed earlier the security processor of this patent may also be deployed in a flow-through configuration. For such a configuration the security processor may include two sets of media interface controller ports as illustrated by blocks 7201 and 7213. The security processor illustrated in Fig. 72 is very similar to that in Fig. 71; however it has multiple media interface controller ports as against the host or coprocessor interface block like block 7101. The number of ports may depend on the line rate per port and the performance of the security processor. The sum of incoming ports line rate should be matched with the processing performance of the security processor to provide security inspection to substantially the entire incoming traffic. A conscious choice could be made to use a higher line rate sum than the processors capability if it is known that not all the traffic needs to be inspected for security purposes. The decision of the traffic that must be inspected may depend on the connection or the session as programmed in the processor from the central manager. The security processor of Fig. 72 may thus be used to provide flow-through security inspection to the traffic and may be used in a flow-through configuration like that illustrated by Figure 68. A flow through configuration may be created for various types of the systems like a switch or a router line card or a host server adapter or a storage networking line card or adapter or the like. In a flow-through configuration the security processor is directly exposed to the traffic on the network. Thus, the central manager and the rules distribution engine may directly communicate to the control processor and scheduler, block 7203 or block 6809, of the security processor. Security processor of block 6802 is similar to the one illustrated in Fig. 72 without the runtime adaptable processor incorporated in it. One

of the issues in a flow-through configuration that needs to be addressed is the latency introduced in the traffic by the security processor. The network switches or routers for example are very sensitive to latency performance of the system, Hence in such a configuration a deep packet inspection can add significant latency to the detriment of the system performance. Hence, the security processors for flow-through configuration of this invention provide a cut-through logic illustrated by block 6807 that is used to pass the data traffic from the input of the security processor to its output incurring a minimal latency to support the overall system performance needs. The control processor and scheduler block 7203 of Fig. 72 provides the cut-through logic and is not illustrated separately. In a flow through configuration, once a match has been found the security processor may create special control packets internal to the system, where the system's switch processor or a network processor or other processors may interpret these messages and perform appropriate action on the packets that utilize the cut-through mode before those packets are allowed to exit the system. Such a protocol may be a proprietary protocol within a system or may utilize a standard protocol as may be appropriate for the system incorporating a flow-through security configuration.

Figure 73 illustrates another version of the security processor which is very similar to that in Figures 71 and 72, with some additional functionality. The additional functionality is provided by classification/rules engine, block 7313, classification/rules database, block 7314 and the database extension controller block 7315. These blocks are similar to those of Figures 20 and 30 described above. These blocks may be used to provide high performance network layer rules processing using a CAM based architecture. The ternary CAM based database may also be used to provide a fast match to specific fields in a network header to create hash keys for connection identification and connection state retrieval or update. The control processor and scheduler decides which parts of a packet to present to the classification/rules engine depending on the rules that are programmed in it versus those programmed in the content search and rule processing engines. A CAM based architecture typically consumes a lot of power and hence may be limited in its applications except when extremely high speeds may be required at extremely low latencies. The content search and rule processing may be able to provide this functionality at much lower power as well as perform the searches for a much larger rule set compared to that in CAM based architecture. The database extension port, block 7315, may be used to extend the CAM database size using external classification/rules engine.

Fig. 74 illustrates a content search and rules processing engine of blocks 7106, 7206, 7306. The content search and rule processing engine each comprises of interface blocks to

various memories that hold the rule state tables distributed by rules distribution engine of Fig. 67. This engine comprises content fetch block, 7405 which fetches the packet data to be analyzed from the packet buffer block 7104 or equivalent from Figures 72 and 73. The character class look-up, block 7406, accept/state transition look-up, block 7407, next state look-up, block 7408, match/action look-up, block 7409 each perform state data fetches from the appropriate state memories of Fig. 71, 72 and 73. The content search state machine, block 7404, includes the state machine used to analyze the fetched character or characters with the data in various tables. The state machine uses the fetched character to index into the character class table to retrieve a column address for the DFA state machine. In parallel the state machine fetches the state transition table data using the current state as an index to retrieve a row address for the DFA state machine. The current state may be initialized to the start state when beginning a new search, otherwise the next state that is retrieved next becomes the current state for the next iteration of the state machine. The row address and the column address are then used to retrieve the next state for the state machine. The retrieved next state index is also used to fetch an action tag if this is a terminating state. An accept state look-up performed in parallel is used to identify if the retrieved state is a terminating state or an error state or a continuing or accepting state. The content search state machine, block 7404 effectively iterates through steps as outlined above until an error is found or a match is found or the packet is exhausted or a combination of these. For connection based look-up functionality, the current internal state like the address pointers, current state and next state and the like are provided to the control processor and scheduler block 7103 or 7203 or 7303 for it to maintain the state in the connection database. When a new packet for a given connection is scheduled the stored internal state of the content search state machine is retrieved and provided to block 7404 to start processing the new packet for the connection as if it has been a continuous stream with previous packets for the connection. Security processors that include a runtime adaptable processor may also comprise a RAP command controller, block 7403, which is coupled to the adaptation controller block 7205 to adapt the runtime adaptable processor to provide the service as needed by the match and the action tag found with that. The action tag may also be provided to the control processor and scheduler for it to schedule the analyzed packet to the runtime adaptable processor. The adaptation controller may use the command(s) provided by block 7403, as a hint or command to get the processor ready with the needed avatar configuration information, if it is not already present as one of the avatars in the runtime adaptable processor, block 7102.

As described earlier the security processor of this invention may be embedded in systems with many different configurations, dependent on the system environment, system

functionality, system design or other considerations. This patent illustrates three such configurations in Fig. 68, 69 and 70. As discussed above Fig. 68 illustrates the security processor in a network line card or an adapter providing flow-through security. In this configuration the security processor may reside next to the media interface as illustrated or
5 after block 6803 closer to the host or back plane interface block 6804. Such decisions are system design decisions and are not precluded from the usage of the security processor of this patent. In a scenario where the security processor incorporates TCP/IP or protocol processing capability, the block 6803 may not be required in some systems. Figure 69 illustrates a look-aside security configuration for a network line card or an adapter. In such a
10 configuration, there exists a master processor which may be a switch processor, network processor, forwarding engine, or classification engine or other processor illustrated by block 6903. The master processor communicates with the central manager of Fig. 66 as described earlier to receive the rules and to provide events back to the central manager, working with the security processor. The master processor may also incorporate functions
15 illustrated by blocks 6902 and 6904. The master processor could also be a TCP/IP processor or other IP processor variations that are feasible from the processors of this patent as well.

Figure 70 illustrates a security and content search acceleration adapter. Such an adapter may be inserted as an accelerator card in multitude of networked systems discussed above
20 like a server, a router, a switch and the like. The security processor on this accelerator card may be coupled to the host bus or back plane directly or through a bridge device like that illustrated by block 7003. The security processor communicates with the host processor or a master processor of the system to receive the packets or content to be inspected and provides the results back. A driver on the host or master processor may perform this
25 communication with the security processor. Such a driver or other software running on the host or the master processor may communicate with the central manager to receive the rules database, or updates to it or provide match results to the central manager based on the actions programmed. The accelerator card may include other devices like a ternary CAM based search engine that may be used to perform network layer security function or
30 connection ID detection or hash key generation or other functions or a combination thereof which may assist to perform network layer and application layer security acceleration functions discussed above.

The security processor of Fig. 71, 72 or 73 may also be used to perform content searches on documents or digital information and be used to create indexes that may be used for
35 accelerated searches like web search capability provided by Google or their competitors.

Using security processor of this invention for such a task can provide significant performance improvement to indexing and searches compared to that done using a general purpose processor based software. For such an application the control processor and scheduler of the security processor may utilize the content search and rules processing engines to
5 perform key phrase searches in data presented to it and get the match indexes. These results can then be used to create a master search index by a process that may run on the control processor and scheduler or another processor of the system that is servicing the content search request from end users. This master index may then be referred to provide quick and comprehensive search results.

10 The security processor of Fig. 71, 72, 73 described above may be coupled with elements of the processor of figures 16, 17, 61 or 62 to provide security capabilities to different versions of protocol processing architectures of this patent. For example, block 6615 illustrates one such variation where the TCP/IP protocol processor is coupled with the processor of figures 71, 72, or 73 to create another security processor with TCP/IP processing. Similar versions
15 may be created by including IP storage protocol processing capability with the security processor or coupling TCP/IP processor with RDMA capability with the security processor of Figures 71, 72 or 73 or a combination thereof. The security processor of Fig. 71, 72 or 73 may also be used in place of the classification engine, block 1703, shown in more detail in Figs. 20 and 30 as described above when the security processor is programmed to search
20 for the classification fields used in block 1703.

Fig. 75 illustrates an example of regular expression rules. As discussed earlier REs have been used since mid 1950s and are used by many popular applications. The figure illustrates a small set of regular expression rules also called a rule tree in this discussion, which can be used to analyze and parse some of the constructs used in HTTP. RE rules like
25 those illustrated with a full set of rules for the language or protocol or the like can be provided to tools like lex, flex and the like which can be used to create lexical analyzers which would then be used to analyze content against the specific language or the protocol. The rules are labeled R1 through R9 in this figure for ease of discussion and don't represent the syntax for a specific tool. The rules illustrate four rule groups: <COMMENT>, <FNT>, <S_TAG> and <INITIAL>. Rules R1 through R3 belong to <COMMENT> rule group and
30 similarly other rules belong to other rule groups as illustrated. Each rule comprises a regular expression and one or more action/tag which can be processed if the regular expression rule is triggered or matched. For example, R9 is part of the <INITIAL> group which searches for an angled bracket ('<') in the content. If the angled bracket is found in the content, then the
35 action/tag is processed. In this rule the action comprises of starting the S_TAG rule group.

The assumption in this illustration is that a rule group labeled <INITIAL> is activated first when the content search starts and then the actions/tags drive the flow to move to activate other rule groups based on the content. In this illustration, when the first angled bracket is detected the rules that are part of S_TAG group are activated and INITIAL group is put on hold. As different content gets processed, the rules flow through the rule tree with the <INITIAL> rule group as the root of the rule tree where the processing of the content search begins. This is an example of the prior art of creating content search rules.

As described earlier, regular expression can be represented using FSA like NFA or DFA. Fig. 76a illustrates Thompson's construction for the regular expression $(xy + y)^*yx$. Thompson's construction proceeds in a step by step manner where each step introduces two new states, so the resulting NFA has at most twice as many states as the symbols and operators in the regular expression. An FSA is comprised of states, state transitions, and symbols that cause the FSA to transition from one state to another. An FSA comprises at least one start state, and at least one accept state where the start state is where the FSA evaluation begins and the accept state is a state which is reached when the FSA recognizes a string. Block 7601 represent the start state of the FSA, while block 7605 is an accept state. Block 7602 represents state 2 and 7604 represents state 3. The transition from state 2 to state 3 is triggered on the symbol x, 7603 and is represented as a directed edge between the two states. Thompson's NFA comprises of 'ε' transitions, 7616, which are transitions among states which may be taken without any input symbol.

Fig. 76b illustrates Berry-Sethi NFA for the regular expression $(xy + y)^*yx$. Berry and Sethi described an algorithm of converting regular expressions into FSA using a technique called 'marking' of a regular expression. It results in an NFA which has a characteristic that all transitions into any state are from the same symbol. For example, all transitions into state 1, 7607, are from symbol 'x'. The other characteristic of the Berry-Sethi construct is that number of NFA states are the same as the number of symbols in the regular expression and one start state. In this type of construction, each occurrence of a symbol is treated as a new symbol. The construction converts the regular expression $(xy + y)^*yx$ to a marked expression $(x_1y_2 + y_3)^*y_4x_5$ where each x_i leads to the same state, 7607. The figure does not illustrate the markings. Once the FSA is constructed the markings are removed. The Fig. 76b illustrates the NFA with the markings removed. As can be seen from the figure, in Berry-Sethi construction all incoming transitions into a state are all dependent on the same symbol. Similarly, a duality of Berry-Sethi construct also has been studied and documented in the literature as discussed earlier, where instead of all incoming transitions being dependent on the same symbol, all outgoing transitions from a state are dependent on the

same symbol. The Berry-Sethi construct is also called a left-biased type of construct, where as its dual is called a right-biased construct.

Fig. 76c illustrates a DFA for the same regular expression $(xy+y)^*yx$. DFA is deterministic in that only one of its states is active at a given time, and only one transition is taken
 5 dependent on the input symbol. Whereas in an NFA, multiple states can be active at the same time and transitions can be taken from one state to multiple states based on one input symbol. There are well known algorithms in the literature, like subset construction, to convert a RE or NFA to a DFA. One point to note for the DFA that is illustrated for the regular expression is that it has fewer states than both the Thompson NFA as well as Berry-
 10 Sethi NFA. The upper bound on the number of states for an N character DFA is 2^N , however expressions that result in the upper bound in the number of DFA states do not occur frequently in lexical analysis applications as noted by Aho, Sethi and Ullman in section 3.7 of their book on Compilers referenced above.

Fig. 77a illustrates a left-biased NFA and its state transition table (prior art). The illustration
 15 is a generic four state Berry-Sethi like NFA with all transitions from each node to the other shown with the appropriate symbol that the transition depends on. For example, state A, 7701 has all incoming transitions dependent on symbol 'a' as illustrated by example transitions labeled 7702 and 7703. When the FSA is in State A, 7701, an input symbol 'd', transitions the FSA to state D with the transition, 7704, from state A to state D. The table in
 20 the figure illustrates the same FSA using a state transition table. The column 'PS', 7711, is the present state of the FSA, while the row 'sym', 7712, is a list of all the symbols that the state transitions depend on. The table 7713, illustrates the next state (NS) that the FSA will transition to from the present state (PS) when an input symbol from those in the sym header row is received. In this FSA, state 'A' is the start state and state C is an accept state.
 25 Hence, if the FSA is in the present state 'A' and an input symbol 'b' is received, the FSA transitions to the next state 'B'. So when the next input symbol is received, the FSA is in present state 'B' and is evaluated for state transition with the row corresponding to present state 'B'.

Fig. 77b illustrates a right-biased NFA and its state transition table (prior art). The illustration
 30 is a generic four state dual of Berry-Sethi NFA with all transitions from each node to the other shown with the appropriate symbol that the transition depends on. For example, state 'A', 7705 has all outgoing transitions dependent on symbol 'a' as illustrated by example transitions labeled 7708 and 7709 where as unlike the left-biased NFA described above, each incoming transition is not on the same symbol, for example transitions labeled 7706
 35 and 7707 depend on symbols 'b' and 'd' respectively. The state transition table in this figure

is similar to the left biased one, except that the FSA transitions to multiple states based on the same input symbol. For example if the FSA is in the present state 'B' and a symbol 'b' is received, then the FSA transitions to all states 'A', 'B', 'C' and 'D'. When an input symbol is received which points the FSA to an empty box, like 7716, the FSA has received a string
 5 which it does not recognize. The FSA can then be initialized to start from the start state again to evaluate the next string and may indicate that the string is not recognized.

The Fig. 77a and Fig. 77b, illustrate generic four state NFAs where all the transitions from each state to the other are shown based on the left-biased or right-biased construct characteristics. However not all four state NFAs would need all the transitions to be present.
 10 Thus if a symbol is received which would require the FSA to transition from the present state to the next state when such transition on the received input symbol is not present, the NFA is said to not recognize the input string. At such time the NFA may be restarted in the start state to recognize the next string. In general, one can use these example four state NFAs to represent any four state RE in a left-biased (LB) or right-biased (RB) form provided there is a
 15 mechanism to enable or disable a given transition based on the resulting four states NFA for the RE.

Fig. 78a illustrates state transition controls for a left-biased and right-biased NFA. The figure illustrates a left-biased NFA with a state 'A', 7800, which has incoming transitions dependent on receiving input Symbol 'S1' from states 'B', 7801, 'C', 7802, and 'D', 7803. However, the
 20 transitions from each of the states 'B', 'C' and 'D' to state 'A', occur only if the appropriate state dependent control is set besides receiving the input symbol 'S1'. The state dependent control for transition from state 'B' to state 'A' is V_2 while those from states 'C' and 'D' to state 'A' is V_3 and V_4 respectively. Transition to the next state 'A' is dependent on present state 'A' through the state dependent control V_1 . Thus transition into a state 'A' occurs depending
 25 on the received input symbol being 'S1' and if the state dependent control for the appropriate transition is set. Thus, one can represent any arbitrary four states NFA by setting or clearing the state dependent control for a specific transition. Thus, if a four states left biased NFA comprises of transition into state 'A', from state 'B' and 'C' but not from the states 'A' or 'D', the state dependent controls can be set as $V_1 = 0$, $V_2 = 1$, $V_3 = 1$ and $V_4 = 0$. Hence if the
 30 NFA is in state 'D' and a symbol 'S1' is received, the NFA will not transition into state 'A', however if the NFA is in state 'B' and a symbol 'S1' is received the NFA will transition into state 'A'.

Similarly, Fig. 78a also illustrates states and transitions for a right-biased NFA. The figure illustrates a right-biased NFA with a state 'A', 7806, which has incoming transitions from
 35 state 'B', 7807, state 'C', 7808, and state 'D', 7809, on receiving input symbols 'S2', 'S3' and

'S4' respectively. However, the transitions from each of the states 'B', 'C' and 'D' to state 'A', occur only if the appropriate state dependent control is set besides receiving the appropriate input symbol. The state dependent control for transition from state 'B' to state 'A' is V_2 while those from states 'C' and 'D' to state 'A' is V_3 and V_4 respectively. Transition to the next state 'A' is dependent on present state 'A' through the state dependent control V_1 . Thus transition into a state 'A' occurs based on the received input symbol and if the state dependent control for the appropriate transition is set. Thus, one can represent any arbitrary four states right-biased NFA by setting or clearing the state dependent control for a specific transition. All state transition controls for a given state form a state dependent vector (SDV), which is comprised of V_1 , V_2 , V_3 , and V_4 for the example in Fig. 78a for the left-biased and the right-biased NFAs.

Fig. 78b illustrates configurable next state table per state. The left-biased state table for 'NS = A', is shown by the table 7811, whereas the right-biased state table for 'NS = A', is shown by the table 7812. The state dependent vector for both left-biased and right-biased NFA state is the same, while the received input symbol that drive the transition are different for the left-biased vs. right-biased NFA states. Thus a state can be represented with properties like left-biased (LB), right-biased (RB), start state, accept state, SDV as well as action that may be taken if this state is reached during the evaluation of input strings to the NFA that comprises this state.

Fig. 79a illustrates state transition logic (STL) for a state. The STL is used to evaluate the next state for a state. The next state computed using the STL for a state depends on the current state of the NFA, the SDV, and the received symbol or symbols for a left-biased NFA and right-biased NFA respectively. The InChar input is evaluated against symbols 'S1' through 'Sn' using the symbol detection logic, block 7900, where 'n' is an integer representing the number of symbols in the RE of the NFA. The choice of 'n' depends on how many states are typically expected for the NFAs of the applications that may use the search processor. Thus, 'n' may be chosen to be 8, 16, 32 or any other integer. The simplest operation for symbol detection may be a compare of the input symbol with 'S1' through 'Sn'. The output of the symbol detection logic is called the received symbol vector (RSV) comprised of individual detection signals 'RS1' through 'RSn'. LB/RB# is a signal (i.e. a left-biased or right-biased signal) that indicates if a left-biased NFA or a right-biased NFA is defined. LB/RB# is also used as an input in evaluating state transition. The STL for a state supports creation of a left-biased as well as right-biased NFA constructs. The LB/RB# signal controls whether the STL is realizing a left-biased or a right-biased construct. The state dependent vector in the form of 'V1' through 'Vn', is also applied as input to the STL.

The SDV enables creation of arbitrary 'n'-state NFAs using STL as a basis for a state logic block illustrated in Fig. 79b. Present states are fed into STL as a current state vector (CSV) comprised of 'Q1' through 'Qn'. STL generates a signal 'N1' which gets updated in the state memory, block 7902, on the next input clock signal. 'N1' is logically represented as $N1 = ((V1 \text{ and } Q1 \text{ and } (LB/RB\# \text{ OR } RS1)) \text{ OR } (V2 \text{ and } Q2 \text{ and } (LB/RB\# \text{ OR } RS2)) \text{ OR } \dots (Vn \text{ and } Qn \text{ and } (LB/RB\# \text{ OR } RSn))) \text{ AND } ((\text{NOT } LB/RB\# \text{ OR } RS1))$. Similar signal for another state 'n', would be generated with similar logic, except that the signal 7901, feeding into the OR gate, 7915, would be 'RSn', which is the output of the 'n'-th symbol detection logic, changing the last term of the node 'N1' logic from $((\text{NOT } LB/RB\# \text{ OR } RS1))$ to $((\text{NOT } LB/RB\# \text{ OR } RSn))$.

The state memory, 7902, can be implemented as a single bit flip-flop in the state logic block discussed below.

Fig. 79b illustrates a state logic block (SLB). The SLB comprises the STL, 7906, Init logic, 7908, state memory, 7910, the accept state detect logic, 7911, the SDV for this state, 7907, start flag, 7909, accept flag, 7912, tag associated with this state, 7919, or action associated with this state, 7913 or a combination of the foregoing. The SLB provides programmable storage for each of the following: SDV, start flag, accept flag, tag field, action field, left-biased or right-biased signal, and the state memory or a combination thereof, which may be programmed using application state controller, block 8000, working with the cluster configuration controller, 8607, and global configuration controller, 8505, discussed below to configure a state of realized FSA in SLB at runtime. The SLB receives current state vector and the received symbol vector which are fed to STL to determine the next state. The realization of a state of an arbitrary NFA can then be done by updating the SDV for the state and selecting the symbols that the NFA detects and takes actions on. Further, each state may get marked as a start state or an accept state or a combination or neither start or accept state through the start and accept flags. The init logic block, 7908, receives control signals that indicate if the state needs to be initialized from the start state or cleared or disabled from updates, or loaded directly with another state value, or may detect a counter value, discussed below, and decide to accept a transition or not and the like. The Init block can be used to override the STL evaluation and set the state memory to active or inactive state.

The STL, 7906, provides functionality as illustrated in Fig. 79a, except that the state memory is included in the SLB as independent functional block, 7910. The state memory, 7910, can be implemented as a single bit flip-flop. When the state memory is set it indicates that the state is active otherwise the state is inactive. The accept detect logic, 7911, detects if this state has been activated and if it is an accept state of the realized NFA. If the state is an accept state, and if this state is reached during the NFA evaluation, then the associated action or the tag is provided as an output of the SLB on the A1 signal, 7916, and an accept

state activation indicated on M1, 7917. The action or the tag could be set up to be output from the SLB on the state activation as well as when the state is not an accept state as required by the implementation of the NFA. This can enable the SLB to be used for tagged NFA implementation where an action or tag can be associated with a given transition into a state.

If there are 'n' states supported in the search engine, discussed later, each SLB needs 'n'-bit SDV which can be stored as a register comprised of flip-flops, 2-bits allocated to start flag and accept flags, 1-bit for LB/RB#, m-bit action storage and t-bit tag storage. Thus if n=16 and m=6 and t=4, then the total storage used per SLB would be a 29-bit register equivalent which is a little less than 4 bytes per state.

Fig. 80 illustrates an NFA based search engine. The search engine may comprise of 'n' SLBs, SLB1 through SLB n, block 8002, Application state control, block 8000, Application state memory, block 8001, Counters, block 8004, Match logic and accept state detect, block 8005, FSA extension logic, block 8006, Symbol detection array, block 8007 or Pattern substitution control, block 8008 or a combination of the foregoing to generate a runtime adaptable NFA search engine. The number of states included per search engine may be any integer number 'n'. For example 'n' may be 8, 16 or 32 or any other integer. A regular expression with up to 'n' symbols may be fully supported by the search engine. If the number of states resulting from the RE are higher than that supported by a search engine, then it is possible to connect two or more search engines together using the FSA extension logic, block 8006, to handle a larger RE. If the number of states of a RE are smaller than 'n', the search engine allows multiple NFAs per search engine. This can result in efficient utilization of the search engines resources unlike solutions that do not allow multiple REs per search engine. There may be 'M' rules that may be supported per search engine where M can be any integer less than 'n'. For instance if 'M' is 4 and 'n' is 16, it is possible to configure four RE rules in the search engine for simultaneous evaluation. There may be many more rules that may be stored in the application state memory, though 'M' is the number that may be simultaneously configured in the SLBs for evaluation. The total number of rules and application contexts stored in the application state memory is a function of the amount of memory allocated per search engine and is not limited to 'n' or 'M'. The SLBs are the state logic blocks illustrated in Fig. 79b. The number of SLBs that are included in the search engine are the same as the number of states supported by the search engine.

The search engine receives as inputs the FSA control and data buses, 'FSACtrl' and 'FSAdata' respectively, the input symbol on the bus 'InChar' and the substitution control and data on the bus 'SubBus'. The other standard inputs like the power, clock, reset and the like

which may be required are not illustrated in this figure or other figures in this disclosure to not complicate the illustrations. Use of such signals are documented in ample in basic hardware design textbooks and is well understood by any one with ordinary skills in the art. Though the description is with respect to the inputs as illustrated, the information may be

5 received on multiplexed busses or other parallel busses and hence this description is not to be taken as a limitation of the invention. These signals or busses are illustrated for the ease of discussion. The search engine also comprises an application state memory which holds application specific NFA rule context for multiple applications like HTTP, SMTP, XML and the like. The application state memory holds information like SDV, initial state vector (ISV),

10 accept state vector (ASV), LB/RB#, Actions/Tags, number of rules and the like for each application NFA. The application state control, block 8000, is used by the search processor's global configuration controller, described later, to configure the appropriate NFA parameters in the application state memory for each of the supported applications. Then the configuration controller can instruct the application state control to program the SLBs using a

15 particular application NFA for performing search for that application. There may be multiple REs or rules per application in the search engine if the number of states for the REs together is less than that supported by the search engine. The application state control can retrieve the appropriate application context from the application state memory and program each SLB with the appropriate information to implement the NFA programmed for the application.

20 The parameters being configured in each SLB comprises the SDV, start state, accept state, LB/RB#, action or tag. The application state controller also configures the symbol detection array, block 8007, using the symbols for the NFA retrieved from the application state memory. Appropriate counter initialization may also be done by the application state control for the NFA before enabling the search engine to start processing input symbols received on

25 'InChar'. The state outputs of the SLBs are concatenated together to form the CSV. The RSV is the output from the symbol detection array. The RSV and CSV are fed into each SLB. The SLBs then perform the state transition evaluation based on their configuration until an accept state is reached or the NFA reaches a null state which indicates that the input string is the one recognized by the NFA or rejected by the NFA respectively. The symbol

30 detection array, 8007, comprises of an array of computation engines, which in the most simplistic form can be comparators, and memory for symbols being searched. The input symbol received on 'InChar' is compared simultaneously with all the symbols and then the appropriate signals of the RSV are set if the symbol being searched matches the received symbol. The symbol detection array may also interact with the counters block, 8004, to

35 enable rules that detect occurrence of symbols before taking a specific transition. The accept state and the action fields from each SLB is used by the match and accept state detection logic, block 8005, to detect if an NFA has reached an accept state and it then

registers the action for the NFA to provide to the global configuration controller discussed later. The action or tag may provide an indication that a RE rule has matched, but may also provide an action that may be taken on the input stream. For example, the action may indicated that the input stream contains a computer virus being searched and hence the

5 input stream should be quarantined for instance or it may detect that confidential information being searched is part of the stream and hence the transmission of this stream should be stopped and the like. The action or tag may also comprise of a pointer to a code routine that the global controller or the control processor, discussed later, may perform. The code routine may cause the configuration of the search engine to change to search for a different

10 set of rules. For example, if the NFA was part of a rule set <INITIAL> discussed earlier and if it triggers, then the action to be taken may be to activate <S_TAG> set of rules and put the <INITIAL> set on hold for instance. There are innumerable type of tasks that may be accomplished using the invoked code routine as a result of the RE rule being matched as would be obvious to one skilled in the art. Complex set of actions can be taken on the input

15 stream or other control functions may be activated or management code may be activated or the like once an NFA reaches an accept state if the associated action or the tag is programmed for such an action to occur. The indication of a match as well as the rule ID being matched may be output on the M-bus, 8009, by the match logic, 8005, for the cluster controller, discussed later, or the global configuration controller to take appropriate action.

20 The action/tag, 8010, carries the associated action or tag to the cluster controller or the global configuration controller which may follow actions as described above. The match logic is set up to support 'M' simultaneous rule evaluation by the search engine.

There are instances of regular expressions where a certain number of events need to be counted before a state transition occurs. For example a RE rule may be created like

25 [3,5]abc which indicates that the RE recognizes a string that has 3 to 5 leading 'a's followed by 'b' followed by 'c'. If the input string is 'aabc' or 'abc' or 'aaaaaabc', then it would not be recognized by this RE, whereas an input string 'aaaabc' or 'aaabc' or 'aaaaabc' would be recognized by this RE. To enable such REs there are counters provided in the search engine as part of the counters block, 8004. There may be 'L' sets of counters to support

30 multiple rule execution simultaneously. If the some rules do not need to use the counters, then those counters may be allocated for use by other rules. The choice of 'L' may be a function of the search engine resource availability and the need of counters for applications being supported using the search engine. The counter values may be read or written to by the application state controller on the configuration of an application state into the SLBs if the

35 NFA configured in the search engine needs counters. The output from the counters block is provided to the SLBs for them to make decisions on the counter values as necessary. The

counters may also be read and written by the symbol detection array for it to support RE rules that depend on counting symbols occurring for a certain number of times like the one described above. The counter block also holds counters to keep track of the string match beginning and the end for each of the 'M' rules. Begin and end counts indicate to the cluster
5 controller or the control processor the start marker for the string match in the input stream and the end marker in the input stream respectively.

The FSA extension logic, block 8006, is used to extend the number of NFA states that can be supported for a RE. Thus if there are REs that need more states than 'n' supported by an individual search engine, it is possible to connect multiple search engines into groups
10 together to support the RE. The FSA extension logic block provides the input and output signals that enable a search engine to know the state of other search engines of the group. The signals may be fed into the SLBs as control signals, though are not explicitly illustrated in Fig. 79a or Fig.79b or Fig. 80.

The search engine may also comprise of a pattern substitution control, 8008, which may
15 hold pattern(s) to be used for substituting the input string as the search engine progresses through NFA. This can support finite state transducer functionality which is well understood and documented in literature. This block may evaluate the transitions taken through the NFA by keeping track of two consecutive CSVs and making decision about the transitions enabled and the associated transducer symbols to be generated for substituting the received
20 input string.

A search engine with 'n=16' states would require $16 \times 4 = 64$ Bytes of storage for SLBs, plus 16 bytes for 16 symbols, and M-counters where if $M = 4$, then there may be 8 bytes for counters if each counter is 16-bit. Thus total storage required for an application state context, discussed below, is around 90 bytes for $n=16$.

25 Fig. 81 illustrates Application State Memory example configuration. The application state memory, 8100, is the same as the block 8001 in the Fig. 80. Fig. 81 illustrates how that memory may be organized to store application state context for various applications. An application state context comprises the information like the SDV, Start State, Accept State, Symbols for the NFA evaluation, actions and tags for the states and the like. This is the
30 information that the application state control can retrieve from the application state memory and configure the SLBs, Counters, Symbol detection array and the match detection and accept detection logic blocks of the search engine. It may also include the indication whether this particular RE needs grouping with other search engine block to support a larger RE than that supported by the number of states supported by the search engine. Thus the

context carries all the state information that is necessary to evaluate an NFA. The number of application state contexts 'N' that can be stored in the application state memory is purely a function of the amount of application state memory allocated per search engine and the storage required per context. When 'n=16', if the amount of storage allocated per search engine is 512 bytes, then 5 application contexts can be stored in the application state context. A 1KB of memory can allow 10 application contexts to be stored. However, if the number of states per search engine is 8, the number of contexts can be up to two times of those in the case of 16 states.

There are 'N' application state contexts (ASC), 8101, 8102, 8103 stored in the application state memory, 8100. ASC1, 8101, is illustrated to be one of the RE rules for anti-spam application, whereas ASCn, 8103, is illustrated to be one of the RE rules for XML application. Each ASC may be further organized into NFA parameters, 8104, and SLB parameters, 8105. The SLB parameters are state specific parameters like Start Flag, Accept Flag, SDV, Action and the like that are applied to the specific SLB. These could simply be a 4-byte memory location. Similarly the NFA parameters are those that are applicable for the entire NFA, like symbols, number of rules per NFA, substitution table and the like. RB/LB information may be common to the entire NFA or may be stored in a per state basis.

The benefit of ASCs is that when the control processor or global configuration controller or cluster controller, discussed later, need to apply the rules for a different application than that currently configured in the search engine, it can make the transition to rules for the new application rapidly by informing the application state control of each search engine. Thus in an array of search engines which may be used to perform evaluation of a large number of REs in parallel, the switching to an application context can be very rapid, unlike if the configuration information has to be retrieved from a global memory or external system memory and provided to each search engine. In the example of n=16 states, it may be possible to switch the application context for each search engine within 20 to 25 clocks. If there are 1000 search engines being used in parallel to do 1000 RE rule searches, and if the application based context has to be loaded into each of them from a global memory or external memory, it can take potentially 20,000 to 25,000 clock cycles, which can bring down the performance of the search processor significantly compared to using ASCs per search engine as in this invention.

Even though the above discussion is with respect to different applications, it may be possible to use different REs for the same applications in different ASC. This can allow the search processor to be able to support a larger number of REs for an application, though there would be performance impact if all the REs have to be applied to the input stream. In such a

case, the input stream may be presented multiple times to the search engine, each time switching the ASC, which in fact then applies a new RE for the application to the stream. Another, usage model of the ASC can be for the same applications' nested RE rules. For example, as illustrated in Fig. 75 the RE rules are nested rules, where only a limited set is being applied to the input stream in parallel. In such a case, one rule from each level of nesting could be set up in individual ASCs and the action programmed could indicate that when the current rule reaches accept state the context needs to be switched to another ASC. For example, R9, R7 and R4 could each be stored in three independent ASCs with Rule R9 being used as the initial context. Now if the input string matches R9, then the context can be switched to bring in rule R7 which then looks for the 'FONT' string. If R7 reaches an accept state, then the context can be switched to R4 and so on. Thus a nested RE rule set for any given application can be easily supported in the search engine of this invention. Thus one can create a runtime adaptable finite state automaton architecture using various runtime configurable and programmable features of the NFA based search engine described above.

As discussed earlier a regular expression can be converted to NFA and DFA using well known methods. Further, it is also well known that an NFA and a DFA for the same RE are both equivalent. However, there are differences in terms of the performance metrics of an NFA vs. a DFA as discussed earlier. DFAs can result in a large storage requirement for the states in a worst case compared to NFAs. However, in practice such is not the case and many times a DFA may have fewer states than a NFA. Thus a content search processor using only NFA based search engines may not realize the most efficient implementation. One can use a combination of NFAs and DFAs to evaluate a set of regular expressions. Further, if the DFA implementation can be more efficient than the NFA, then more rules can be supported on a single integrated search processor compared to the one that purely uses NFAs. DFAs may have transitions to multiple states from a state, each transition taken on a different input symbol. Thus to process a state, one would need to detect multiple symbols and based on the input symbol choose the appropriate transition. In DFAs only one state is active at a time by the deterministic nature of DFAs. Thus, it may be possible to select DFAs that are like the one illustrated in Fig. 76c, where the number of states in the DFA are less than or comparable to those in an NFA and the number of transition from a state are also limited, in this case an upper bound of two and program them for evaluation using a DFA based search engine like that illustrated in Fig. 82. As long as the number of states of the DFA are not more than that can be held in the DFA instruction and state flow memory, described below the DFA could be chosen instead of the NFA for evaluation.

Fig. 82 illustrates a DFA based search engine. The DFA search engine comprises of DFA operations engine, 8204, DFA state controller, 8203, DFA context control, 8200, DFA Instruction and state flow memory, 8201, counters, 8202, DFA extension logic, 8208, occurrence counter, 8206, Min & Max counters, 8207, or pattern substitution control, 8205, or a combination of the foregoing to generate a runtime adaptable DFA search engine. The DFA search engine also has the facility to create extension of the FSA using the DFA extension logic, 8208, which provides functionality similar to the FSA extension logic, 8006, of Fig. 80. Thus multiple DFA blocks can be grouped together to support REs that result in larger DFAs than those supported by one DFA search engine. The DFA operations engine is an engine which is a small processor that supports executing targeted DFA specific operations. An example of the type of DFA operations are illustrated in Fig. 83. DFA State Controller, 8203, receives commands from the cluster controller or the global configuration controller to program the instruction memory & state flow memory, 8201 with the instruction sequence using the DFA operations that would form the DFA flow. The DFA context Controller, 8200, controls the application and the DFA context that is currently being executed. The DFA search engine also supports application state contexts similar to the NFA search engine of Fig.80 and they are stored in the instruction and state flow memory, 8201. The application context for the DFA is primarily comprised of instruction sequence for the RE rules for each application. Similar to the NFA search engine discussion above, the DFA search engine also can hold multiple REs for the same application and thereby support the RE nesting as described earlier for NFAs. The DFA operations engine, 8204, is essentially a scalar processing engine that supports all DFA operations, like the example described below. DFA operations engine may be a multi-way superscalar architecture where multiple instructions can be processed simultaneously. The multi-way instructions would essentially be state transition evaluation instructions which would take the DFA to one of the multiple next states based on the received input symbol. If the DFA operation engine supports two simultaneous evaluations to support DFAs with at most two transitions per state, then each RE could be converted into DFAs to meet such requirement or one could compare the cost of implementing an RE as an NFA or a DFA and choose the least costly realization by mapping that RE to the appropriate search engine on the content search processor. The DFA operations engine in general does not retain state information, since the DFA instruction flow holds the knowledge of the state. Further, the DFA operations engine evaluates each incoming symbol through the same engine and hence it does not have the hardware overhead of multiple symbol detection array unlike the NFA search engine. This causes the DFA search engine to be significantly cheaper to implement compared to an NFA search engine where the number of device utilization difference could be up to 5 times. DFA context controller, 8200, provides the context and the pointer of the

instruction within the context to the instruction memory, which retrieves the instruction vector(s) to be evaluated by the DFA operations engine, 8204. The instruction vector is comprised of the instruction or the operation, the expected symbol(s), next state pointer(s), a default state, or control information or a combination of the foregoing. The instruction vector
5 may contain information for multiple transitions if a multi-way architecture is selected for the DFA search engine. The DFA operations engine executes or evaluates the instruction against the received/input symbol and makes a decision as to which state to transition to next and start its evaluation. If there is no match with the received symbol and the expected DFA operation, the flow transitions to a default state which may be the start state or an error
10 state. If the DFA operation engine reaches an accept state then an appropriate match signal is generated on the M-bus by the DFA operations engine along with the tag or the action associated with this accept state. This information is processed by the cluster controller or the global configuration controller, who decide the next course of action to follow for the input. The DFA operations engine has access to occurrence counter, 8206, min & max
15 counters, 8207 and character counters, 8202 which are used to support various DFA operations. The character counters are used to mark beginning and end of a string that matches the DFA rule. Pattern substitution control, block 8205 and DFA extension logic, block 8208, provide functionality similar to that discussed for the blocks 8008 and 8006 respectively, illustrated in fig. 80.

20 Fig. 83 illustrates example DFA operations. The operations have been defined to model common regular expression operators. The number and type of DFA operations can be different from those illustrated without deviating from the teachings of this patent. For example an illustrated operation 'EQ X' may be used to evaluate if the received symbol is equal to 'X'. If the evaluation results in an affirmative answer the transition that requires this
25 evaluation can be taken. The DFA operations illustrated form the instructions that get processed or evaluated by the DFA operations engine, 8204. DFA instructions comprise more fields than those illustrated in the Fig. 83. The DFA instruction may comprise DFA Operation, like that illustrated in Fig. 83, expected symbol(s), target instruction address (s) to branch to on successful evaluation of the instruction, alternate target instruction address if
30 the evaluation is not successful, or other control fields like setting of counters, retrieval of counters, output an accept state flag, action or tag and the like or a combination of the foregoing. The number of symbols being evaluated and the target instruction addresses would depend on the number of parallel transitions being supported. When the DFA operations engine is two-way superscalar the number of symbols would be two and there
35 would be two target addresses associated with those beside the default address. Each instruction of the DFA operation would represent the evaluation of transitions from a state of

the DFA. Thus when a branch is taken to a target instruction address the instruction at the target would typically be the one for the state that the DFA transitions into, unless this is an error state in which case the DFA may be restarted at the initial or the start state. For example one could code the DFA illustrated in Fig. 76c as follows using some of the DFA operations illustrated in Fig. 83. The example below is not limiting format for the DFA operations and is used for illustration purpose to explain the DFA search engine operation better.

State/IP	Op1	S1	T1	Op2	S2	T2	Other
A: 0	EQ	'x'	1	EQ	'y'	2	0, start count
B: 1	EQ	'y'	0	EQ	-	-	0
C: 2	EQ	'y'	2	EQ	'x'	3	0
D: 3	GEN_FLAG	-	-	EQ	'y'	0	0, action, tag, flag, end count

The fields of the instructions above are op1 (first operation), S1 (symbol1), T1 (IP Target1), op2 (second operation), S2 (symbol2), T2 (IP Target2) and other fields (like Default target, init start count, init end count, action, tag and the like). For evaluation of the DFA illustrated in Fig. 76c, such a code may be stored in the DFA instruction and state flow memory from where each instruction is retrieved and processed. The column with the header "State/IP" is representing the corresponding state of the DFA and the instruction pointer (IP) value for the corresponding instruction. There may be a preamble code to that illustrated that may need to be executed before the DFA code to setup appropriate counters, error state processing and the like. However, in the illustration of the example above the start state 'A' and its IP is used as the default target if the input symbols are not those expected by the DFA. The DFA evaluation starts in the start state at IP=0. The instruction at IP=0 is first fetched and evaluated against the input symbol. If the input symbol is a 'y' then the DFA is expected to transition from state 'A', 7612, to state 'C', 7614, however if the received symbol is 'x' then the transition is to state 'B', 7613. Assuming that the input symbol is a 'y', then the instruction evaluation of the first instruction results in an affirmation of the comparison of symbol2 (S2) with 'y' and hence the associated target IP of 2 is selected as the next instruction to fetch and evaluate against the next symbol. The IP=2, corresponds to state 'C', which is expected to be the transition that is defined by the DFA of fig. 76c. The DFA search engine now processes instruction at IP=2 and then follows the appropriate flow as defined by the instructions. When the execution reaches IP=3, which corresponds to the accept state 'D', 7615, the DFA search engine outputs a match flag and associated tag or action to indicate that a string matching the current DFA has been found. The DFA engine also

records the end count to indicate where the string match ends in the input stream. The DFA engine with multi-way execution like the one illustrated above may also provide priority to the multiple operations being performed simultaneously, where the first operation's affirmative result may get a higher priority over the other operations. In general, since the DFAs are
5 deterministic, only one of the listed operations of the instruction should result in affirmative result though such a mechanism may be created to avoid execution ambiguity.

Though the example above used only two types of operations, there are many other operations that DFA operations engine can perform as illustrated in Fig. 83. There may be RANGE operation defined which enables one to search for a symbol 'S' to fall within the
10 range of ASCII codes for two symbols with 'X' indicating the lower bound of the range and 'Y' indicating the upper bound of the range. 'X' and 'Y' in this case may be any 8-bit or 16-bit extended ASCII codes. Similarly, a set of operations are defined that enables the DFA code to loop in a given state until a certain condition is met or if a certain number of occurrences of a certain symbol occurs. There are other ways to realize these operations which would be
15 within the scope of this invention as may be obvious to one skilled in the art.

The DFA based search engine of Fig. 82 is essentially an application specific processor which operates on targeted instructions like those illustrated in Fig. 83. The DFA search engine operates by having the DFA context control, 8200, fetch the instructions from the DFA Instruction and state flow memory, 8201, which are then evaluated by DFA operations
20 engine, 8204, by performing the operation indicated by the instruction, which may typically involve evaluating the input symbol received on 'InChar' with those expected by the instruction and then decide the next instruction to fetch based on the result of the evaluation. The DFA operations engine, 8204, provides the information about the next instruction to the DFA state and context control blocks 8203 and 8200, which in turn provide the next
25 instruction address to the DFA instruction and state flow memory, 8201. The performance of this processor can reach the same level or higher as that of state of the art microprocessors using advanced process technologies and thereby achieve line rate performance of 1Gbps to 10Gbps and higher. Thus one can create a runtime adaptable finite state automaton architecture using various runtime configurable and programmable features of the DFA
30 based search engine described above.

A DFA search engine as illustrated in Fig. 82 may be implemented in 50 to 200 logic elements in today's FPGA, depending on the supported DFA operations and multi-way execution support. The instruction and flow state memory may be 512 bytes per search engine or less depending on the number of application contexts that get supported and the
35 number of states that need to be supported. This can allow over a thousand DFAs to be

included in a single FPGA which can have close to 200 thousand logic elements. Thus DFA and NFA based content search processor may be implemented in today's best of class FPGAs supporting one to two thousand REs in one FPGA which may be operated at the clock rate achievable in the FPGA. Higher number of REs can be supported in an FPGA format when multiple application rules are taken into account, which can drive the number into multiple thousands of rules. This invention thus supports creating content search solutions using FPGAs beside ASICs, ASSP and other forms of integrated circuits, which may be useful when the time to market is important and the volumes are not high enough to justify the expense of creating ASICs, ASSPs or other ICs. When an NFA and DFA search engines of this invention are implemented on the best of class process technologies like 90nm or 65nm that are used in today's leading edge microprocessors, the performance of these engines can be taken to the same or higher frequencies and thus can enable very high line rate processing of regular expressions unlike today's microprocessors. Unlike below 100 Mbps performance of a 4-GHz processor for evaluating a few hundred NFAs or DFAs as described earlier, a thousand RE content search processor with NFA and DFA search engines of this invention could achieve clock rates similar to the leading microprocessors and evaluate around one input symbol per clock, which may be a single character or more characters, thereby achieving well above 10Gbps performance. Further, with the device densities growing with small geometry process technologies like 90nm, 65nm and lower, the number of REs that can be incorporated on an IC can be in multiple thousands depending on the choice of the mix of NFAs and DFAs. By using a mix of NFAs and DFAs on the content search processor it would be possible to create a series of products that offer different numbers of both in a chip, for example one product can offer 25% DFAs and 75% NFAs, while another could offer 75% DFAs and 25% NFAs for instance. Thus it may become possible to meet varying market needs. If an application requires more REs than those provided on a single content search processor, multiple processors can be used in parallel to increase the number of REs. Thus, it is possible to use this invention to support search applications requiring thousands of REs without the need to grow the memory usage like that would be needed in a composite DFA that make it uneconomical for general use.

Fig. 84 illustrates a runtime adaptable search processor. The processor comprises a control processor and scheduler, 8402, a coprocessor/ host interface controller, 8401, data/packet buffers, 8403, adaptation controller, 8404, a runtime adaptable search processor array, 8405, configuration and global memory, 8406, or memory controller, 8407, or a combination of the foregoing. The runtime adaptable search processor is also referred as a search processor or adaptable search processor or content search processor or runtime adaptable content search processor in this invention. Similarly, runtime adaptable search processor

array is also referred to as the search processor array in this invention. The search processor may be embedded in multiple types of systems in varying configurations similar to those in Fig. 68, Fig. 69 and Fig. 70, by substituting the security processor of these figures with the search processor. Search solutions derived from the processors of this invention

5 comprise line cards which may incorporate the search processor, security processor, SAN protocol processor, TCP/IP processor or runtime adaptable protocol processor or various other processors disclosed in this patent or a combination of the foregoing. The line card configuration and the architecture may vary with the specific system and the application. Three types of line card architectures, a) flow-through b) look-aside and c) accelerator card,

10 are illustrated in this patent to illustrate usage models for the processors of this patent. Fig. 68, Fig. 69 and Fig. 70 illustrate these configurations using a security processor based system, though it could also be based on the search processor or other processors of this patent as discussed earlier. Blocks 6612 and block 6613 illustrate two of these types of card configurations. The security processor illustrated in these cards may be substituted with the

15 search processor disclosed in this patent. There are various different variations of the search processor that can be created depending on the functionality incorporated in the processor. Blocks 6614 and block 6615 illustrate two versions of a security processor which could also be used to illustrate search processors by replacing the security processor with search processor. Block 6614 illustrates a security processor comprising at least a content

20 search and rule processing engine coupled with a runtime adaptable processor. A similar search processor can be created by substituting the runtime adaptable processor with a runtime adaptable search processor array, 8405, as disclosed in this application. This processor is similar to that illustrated in Fig. 71 with the same substitution and is described in detail above within the context of a security processor. Block 6615, illustrates the security

25 processor of block 6614 coupled with a TCP/IP processor or a protocol processor to provide more functionality usable in a security node as a security processor. Similarly, a search processor may also be coupled with a TCP/IP processor or a protocol processor to provide more functionality usable in a content search processor node that may be used in a network as a network element or a client node or a server node. Such a processor may be used to

30 enable application aware networking capabilities to achieve intelligent networking. The choice of the search processor may depend on the system in which it is being deployed, the functionality supported by the system, the solution cost, performance requirement, or other reasons, or a combination thereof. The search processor may use one or more ports to connect to external memories, block 6616, which may be used to store additional rules

35 beyond those stored on the search processor, or other intermediate data or packets or other information as necessary to perform various functions needed for search processing. The memories may be of various types like DRAM, SDRAM, DDR DRAM, DDR II DRAM,

RLDRAM, SRAM, RDRAM, FCRAM, QDR SRAM, DDR SRAM, Magnetic memories, Flash or a combination thereof or future derivatives of such memory technologies. The inventions disclosed in this patent enable many variations of the architectures illustrated and may be appreciated by those skilled in the art that changes in the embodiments may be made
5 without departing from the principles and spirit of the invention.

The search processor of Fig. 84 may be used as a coprocessor in a system which needs acceleration for content search of local content, remote content or content received in network packets or the like. The configuration in which this processor may be used is similar to that illustrated in Fig. 69 and Fig. 70 where the security processor of the blocks 6906 and
10 7005 may be replaced with the search processor of Fig. 84 or other variations of the search processor as described above. The search processor may be incorporated on an add-in card which may act as an acceleration card in a system as illustrated in Fig. 89, where the search processor, 8904, may be that of Fig. 84 or other variations of the search processor as described above. The search processor may also be configured in a form illustrated by Fig.
15 71 where the runtime adaptable processor, block 7102, may be substituted with runtime adaptable search processor array, block 8405. Such search processor can have the benefit of allowing the number of rules the processor can support to grow beyond the number supported by the runtime adaptable search processor array. Thus if this search processor array is able to support a couple thousand REs, then if the user need to grow the number of
20 REs beyond that, they may be supported as a composite DFA that may be implemented on the content search/rule processing engines, 7106. Thus the user may be able to grow the number REs upto a point that the external memories can support without needing to add another search processor. The content search processor of this application may also be used as a security processor in applications that require deep packet inspection or require
25 screening of the application content like anti-spam, anti-virus, XML security, web filtering firewalls, intrusion detection/prevention and the like. Thus all figures in this application that refer to security processor may also incorporate a search processor.

The search processor receives the content it needs to search from the master processor in a coprocessor configuration like that of Fig. 69, Fig. 70, Fig. 71, Fig. 73 or Fig. 89 or it may
30 receive the content from the network when coupled with the network interfaces like those in Fig. 68 or Fig. 72. The control processor and scheduler, 8402, deposits data or information content to be searched into data/packet buffers, 8403, for scheduling them into the search processor array when the processor array is ready to receive it. In general the search processor can meet the network line rates from below 100 Mbps, 1Gbps to 10Gbps and
35 higher, and hence the packets may not need to be buffered. However, it may be necessary

to buffer the packets if the packets or content comes into the search processor at a rate higher than what it can handle. It may also be necessary to schedule the packets that belong to the same flow (as in the same transport layer connection for instance) to be processed by the search processor in the correct sequential order. The packets of such flows may also get stored in the data/packet buffers until packets in the flow before them get processed. The search processor may also need to provide inspection ability across multiple packets in a connection between a source and a destination. The control processor and scheduler, block 8402, provides such functionality as well. The control processor may store the internal processing state of the runtime adaptable search processor array in a connection database which may be maintained in the configuration and the global memory, 8406, or in the off-chip memory. The control processor and scheduler looks up the execution or analysis state for a given connection when a packet corresponding to the connection is presented to it by the master processor or the incoming traffic in a flow-through configuration. The connection ID may be created by the master processor and provided that to the search processor with the packet to be inspected or the search processor may derive the connection association from the header of the packet. The connection ID may be created in the IP protocol case by using a 5-tuple hashing derived from the source address, destination address, source port, destination port and the protocol type. Once the connection ID is created and resolved in case of a hash conflict by the control processor and scheduler, it then retrieves the state associated with that connection and provides the state to the search processor array, block 8405, to start searching from the state of the connection. This mechanism is used to create multi-packet searches per connection and detect any search strings or security violations or threats that span packet boundaries. For example, if there is a rule defined to search for "Million US Dollars" and if this string appears in a connection data transfer in two separate packets where "Million U" appears in one packet and "S Dollars" appears in another packet then if a connection based multi-packet search mechanism of this patent is not present the security violation may not be detected since each packet individually does not match the rule. However, when the multi-packet search is performed, no matter how far apart in time these two packets arrive at the search processor, the state of the search will be maintained from one packet to another for the connection and the strings of two packets will be detected and flagged as a continuous string "Million US Dollars".

The control processor and scheduler, schedules the packets or data to the runtime adaptable search processor array in the appropriate order. The runtime adaptable search processor array, block 8405, adaptation controller, block 8404 and configuration and global memory, block 8406 are similar to those illustrated in figures 62, 63, 64 and 65. The runtime

adaptable search processor array and the associated blocks provide similar functionality with appropriate logic enhancements made to couple to the control processor and scheduler of the search processor. The runtime adaptable search processor array and its components are described below. The search processor array is presented with each character of the incoming packet or data, which it then examines for string match with the RE rules
5 programmed in them. When a string is matched the search processor array provides that indication along with other information like the RE rule identification, the action to be taken, the tag associated with this rule, the start count of the match, the end count of the match and the like to the control processor and scheduler, which may then take appropriate action
10 depending on the rule or rules that have been matched. The controllers inside the search processor array may also take appropriate action(s) as required as discussed below.

Fig. 97 illustrates search compiler flow which is used for full and incremental rules distribution. The search compiler of Fig. 97 allows an IT manager to create search and security rules of different types as discussed above in the discussion of Fig. 67 and enable
15 them to create a layered and/or pervasive security model. The compiler flow would be provided with the characteristics of the specific nodes like the security capability presence, the rules communication method, the size of the rule base supported, the performance metrics of the node, deployment location e.g. LAN or SAN or WAN or other, or the like. The compiler flow then uses this knowledge to compile node specific rules from the rule set(s)
20 created by the IT manager. The compiler comprises a rules parser, block 9704, for parsing the rules to be presented to the FSA Compiler Flow, block 9706, which analyzes the rules and creates rules database used for analyzing the content. The rule parser may read the rules from files of rules or directly from the command line or a combination depending on the output of the rule engines. The rules for a specific node are parsed to recognize the
25 language specific tokens used to describe the rules or regular expression tokens and outputs regular expression rules, 9705. The parser then presents the REs to the FSA compiler flow, described below in detail for Fig. 96. The FSA compiler flow processes the incoming REs and generates NFA and DFA for the RE. It then decides whether the RE will be processed as a DFA or an NFA based on the cost of using one vs. the other and also the
30 search processor capability for the node. It then generates the NFA or DFA rule in a format loadable into the search processor and stores it in the compiled rules database storage, 9708.

Rules distribution engine, block 9709, follows the central manager and rules distribution flow illustrated in Fig. 57 and Fig. 58 with appropriate changes to communicate to the right
35 configuration of the search processor. The search rules may be distributed to the host

processor or a control plane processor as illustrated in Fig 58 or to the control processor and scheduler, block 8402, or a combination thereof as appropriate depending on the node capability. The rules may be distributed using a secure link or insecure link using proprietary or standard protocols as appropriate per the specific node's capability over a network.

- 5 The control processor and scheduler, block 8402, communicates with the rules distribution engine, block 9709 to receive appropriate compiled rule tables prior to starting the content inspection. It programs the received rules into the appropriate NFA or DFA search engines, described earlier, working with the adaptation controller, 8404. There may be multiple rules being stored in each search engine dependent on the number of application contexts
- 10 supported by the search processor. Once the rules distribution engine provides the compiled rules to the control processor and scheduler and they are setup in their respective engines, the search processor is ready to start processing the data stream to perform content inspection. The search processor state configuration information is received via the coprocessor/host interface controller or a media interface controller, not illustrated. The
- 15 search processor of this patent may be deployed in various configurations like a look-aside configuration illustrated in Fig. 69 or flow-through configuration illustrated in Fig. 68 or an accelerator adapter configuration illustrated in Fig. 70 as well others not illustrated which can be appreciated by persons skilled in the art. In a look-aside or an accelerator adapter configuration, the search processor of this patent is under control of a master processor
- 20 which may be a network processor or a switch processor or a TCP/IP processor or classification processor or forwarding processor or a host processor or the like depending on the system in which such a card would reside. The control processor and scheduler receives the configuration information under the control of such master processor that communicates with the rule engine to receive packets that contain the configuration
- 25 information and passes it on to the search processor. Once the configuration is done, the master processor provides packets or data files or content to the search processor for which content inspection needs to be performed using the coprocessor or host interface. The coprocessor or the host interface may be standard buses like PCI, PCI-X, PCI express, RapidIO, HyperTransport or LA-1 or SRAM memory interface or the like or a proprietary bus.
- 30 The bandwidth on the bus should be sufficient to keep the content search engine operating at its peak line rate. The search processor may be a memory mapped or an IO mapped device in the master processor space for it to receive the content and other configuration information in a look-aside or accelerator configuration. The search processor may be polled by the master processor or may provide a doorbell or interrupt mechanism to the
- 35 master to indicate when it is done with a given packet or content or when it finds a match to

the programmed rules. The control processor and scheduler, block 8402 and the interface controller, block 8401 work with the master processor to provide the above functionality.

The control processor and scheduler stores incoming packets to the packet buffer, block 8403, and schedules the packets for processing by the search processor array, block 8405, as they become available to analyze the content. The scheduler maintains the record of the packets being processed by the specific engines and once the packets are processed it informs the master processor. The search processor array informs the control processor and the scheduler when it has found a match to a rule and the action associated with that rule. This information may in turn be sent by the control processor to the master processor, where the master processor can take specific action indicated by the rule for the packet. The actions may be one from a multitude of actions like dropping the packet or dropping a connection or informing the IT manager, or the like, as discussed earlier.

Fig. 85 illustrates runtime adaptable search processor array. The runtime adaptable search processor array is similar to the runtime adaptable processor illustrated in Fig 64 and discussed above. The functionality of various blocks of the RAP and search processor array is similar. The search processor array illustrates that each compute cluster, 6401 (1) through 6401 (Z), in Fig. 64 may be either a compute cluster or a search cluster or a combination, as illustrated by blocks 8501 (1) through 8501 (Z). If the cluster is a compute cluster, then its functionality is the same as that discussed in Fig. 64 and Fig. 65, however if the cluster is a search cluster than its functionality is different as illustrated in Fig. 85 and discussed below. The runtime adaptable search processor array could be used instead of runtime adaptable processor in Fig. 62 and Fig. 63 and offer the runtime adaptable search processing and runtime adaptable compute processing capability to the adaptable TCP/IP processor for capabilities discussed above as well as deep packet search, application layer search, upper layer (layers 4-7) security and the like that can utilize search for a large number of rules simultaneously. The global memory and global memory controller, 8502, provides functions similar to those described for block 6402 above, and may also interact with the global and configuration memory, 8406, of the runtime adaptable search processor illustrated in Fig. 84.

Fig. 86 illustrates an example search cluster. The search cluster is similar to the compute cluster illustrated in Fig. 65, however, instead of compute elements, the search cluster comprises of search engines. A search engine may be DFA search engine, illustrated in Fig. 82 or a NFA search engine, illustrated in Fig. 80. The search cluster may comprise of a combination of DFA based search engines and NFA based search engines, or it may comprise of only one type of search engines, either DFA based or NFA based. The

selection may be made at the time of implementation of the search processor for various design optimizations like keep a regular structure within a search cluster which may result in efficient area utilization, power grids, clock network and the like. It would be obvious to those skilled in the art that various configurations of the search clusters are feasible from the teachings of this patent and it is not feasible to describe them all. The search cluster may comprise of search engines, cluster routing switch, 8605, cluster memory and flow controller, 8606, or cluster configuration controller, 8607 or a combination of the foregoing. The block cluster memory and flow controller, 8606, comprises a cluster memory, a cluster memory controller and a cluster flow controller. The search processor may comprise some clusters, that may be DFA based, while others may be NFA based or a combination of DFA and NFA based search clusters. Thus it is possible to create a spectrum of products that may offer varying numbers of DFA based and NFA based clusters as described above. The search engines, 8601(1) through 8601(Z) may be interconnected with each other as necessary to create RE rules that may be larger than those supported by single search engine (SE). It may also be possible to create groups of SEs that all are part of a RE rule group like those illustrated in Fig. 75. SEs that belong to a group may be enabled or disabled by the controllers, 8707 and 8709 depending on the results of the search by the currently enabled group of SEs. The content that needs to be searched for the rules configured in the search engines is presented to the search engines by the search processor array, Fig. 85, through cluster configuration controller, 8607, the cluster routing switch, 8605, and the inter-cluster routing switch, 8503. The results of the search may also be carried through the similar route from the search clusters to the control processor and scheduler, 8402. The search results are also communicated to the cluster memory and flow controller, 8606 and the cluster configuration controller, 8607 which may be configured to take actions as required by the rules that get matched. These controllers may inform the actions that need to be taken to other clusters such that all the search clusters are synchronized with each other in terms of the rule groups that are active as well as the current application context that may need to be active for each cluster.

The search engines in a search cluster are configured by the cluster configuration controller, 8607, which interacts with the global configuration controller, 8505. The global configuration controller is used for global configuration control and interacts with the adaptation controller, 8404, control processor and scheduler, 8402, configuration and global memory, 8406 and the memory controller, 8407 to receive, store and retrieve various hardware configuration information as required by the rules being configured by the user or an IT manager using the search compiler flow illustrated in Fig. 97. The search compiler compiles the rules as per the capabilities of a specific search processor of search device and then interacts with the

control processor and scheduler, 8402, through a host or master processor or otherwise as described earlier, to configure the search processor array with these rules. These rules are communicated by the control processor, interacting with various blocks listed above, to the cluster configuration controller. The cluster configuration controller then configures the

5 application state memory, 8001, for an NFA based search engine and DFA Instruction and State flow memory, 8201, for a DFA based search engine. Once all the search engines in all the clusters are configured with appropriate rules and the application contexts the content search may start. As described earlier there may be a need for the search processor to maintain the information of flows or sessions to continue searching of content inside packets

10 of the same flow that may arrive at the search processor at different times. Similarly in solutions that require searching of messages or files or other content which may be sent to the search processor by a host processor, there may be a need to maintain the context if the content is sent in chunks of bytes or pages or segments or the like. The flow context may be maintained in the global memory, 8406, of the search processor or may be stored in the

15 memory coupled to the memory controller, 8407. When the number of search engines is large, the amount of information that may need to be stored for a given context can cause a performance issue. The DFA and NFA search engines of this patent may need to save a minimum of one 32-bit word per context. Hence, for example if the number of search engines are 1024, distributed amongst 64 clusters, with 16 search engines each then one

20 would need to store and retrieve up to 1024 32-bit words from global memory. This can be a significant performance barrier for the search processor that may need to support 1Gbps to 10Gbps line rate. Hence, for applications that need high line rate performance, this invention enables the clusters to locally store the flow context information in the cluster memory, 8606. For the current example, if the search engines are organized in a 4 X 4

25 array, and there are four ports into the local memory, then on a per cluster basis four 32-bit words of flow context need to be stored and retrieved for the flow context switch. Since all the clusters can access their local memories in parallel, the flow context switch can thus be accomplished in 8 clock cycles compared to 2048 clock cycles when the flow information has to be stored and retrieved from the global memory. The loading of a new context would

30 require 4 clock cycles, which would be the minimum time needed to switch the context where the storing of the context being swapped out could be done in background while the search of the new context begins. Thus, the patent of this invention can solve a major performance bottleneck that may exist in architectures that require the context to be stored in the global memory.

35 Fig. 87 illustrates intra-search cluster rule groups and switching example. As discussed earlier, regular expression rules may be grouped together as illustrated in Fig. 75. The

search processor of this invention allows the grouping of search engines to enable the rule groups of regular expressions. Fig. 87 illustrates two such rule groups which may be created as per the application need and the application context. Rule group 1, 8701, is illustrated to comprise of three search engines, while Rule group 2, 8702, is illustrated to comprise at least 8 search engines. The number of search engines in a rule group may vary based on the application context and a rule group may go across one or more search clusters. Thus a very flexible search processor can be created by this invention. The illustration primarily shows the rule groups confined to a search cluster. The figure illustrates an example of switching of rule groups. Assuming search engine group 1 is enabled to search for content first and the rule group 2 is disabled. Once the content search begins, if one of the rules in rule group 1 finds a match, it can inform the cluster memory and flow controller, 8707 and the cluster configuration controller, 8709, of the action that may need to be taken as a result of the match as illustrated by sending a message S1, 8703 to the memory controller which in turn may send a similar message Sn, 8711, to the configuration controller, 8709. The message Sn is used to illustrate multiple interactions between the controllers to avoid complicating the figure for each step of the message flow. The controllers in turn send a message S2, 8704, to the rule group 1 to deactivate the search engines in the group and send message S3, 8705 & 8706, to rule group 2 to activate the search engines in the rule group 2. Once the rule group switching occurs, the newly activated rule group 2 starts inspecting the content for its rules, while rule group 1 stops inspecting the content. Thus nested rule groups may be created using this invention. The cluster configuration controller, 8709, may interact with global configuration controller, 8505 to make a switch to another rule group that may reside in another search cluster(s) or to make a switch to a different application context across the search processor as may be required by the search result.

Fig. 88 illustrates example rules configuration and context switch. The figure illustrates the runtime adaptable search processor array of Fig. 85, with different type of rules configured in each of the search clusters. The blocks 8801 (1) through 8801 (Z) are the search clusters similar to the blocks 8501 (1) through 8501 (Z). The figure illustrates two contexts for each cluster. The search engines in search cluster 8801 (1) is configured with HTTP and SMTP rules beside any others not shown, whereas 8801 (L) is configured with HTML and SMTP. The label "I: <AppName1>" is meant to indicate that the cluster has <AppName1> application configured to be the Initial configuration. The label "S: <AppName2>" is meant to indicate that the cluster switches to AppName2 under the control of the global configuration controller. There are some clusters marked with label "I: WAIT", which is meant to illustrate that those clusters are not active during the Initial phase of the search. However, those

clusters may switch to the application context shown with the "S: <appName>" label. When the search processor is configured like the illustration in Fig. 88 with the initial set of clusters being active, then those clusters initially inspect the incoming content. In this example each cluster is searching for a particular type of application content being received, which may be done in a networked system for instance. Assuming that the incoming packet belongs to an SMTP application, then the search engines in search cluster, 8801 (3), detect the SMTP packet and inform the cluster controller, which then interacts with the global configuration controller by sending a message, 8817, which indicates that SMTP content is found and hence other clusters should apply rules for SMTP which may be anti-spam rules or parsing the SMTP content into its components or the like. The global configuration controller then sends the messages, 8818, 8819, 8820, to all the clusters for them to switch their context to SMTP if they have that context configured and start applying those rules to the content being received. Thus, the search processor array can be adapted at runtime to perform complex search operations. Though the examples of Fig. 87 and Fig. 88 are fairly simple to keep the description simple, there can be a lot more complex scenarios that can be created using the runtime adaptable search processor as may be appreciated by those skilled in the art.

Fig. 89 illustrates a computing device with content search accelerator. The computing device may be a server, a workstation, a personal computer, a networking device like a switch or a router or other type of computing device. This is one type of configuration in which a content search accelerator using one version of the content search processor of this invention may be used. The figure illustrates a computing device comprising one or more CPUs, 8900 (1) through 8900 (n), at least one chipset, 8902, at least one memory component, 8901, with at least one content search accelerator, 8903, and zero or more adapters providing other functions. The content search accelerator may comprise of search processor, 8904. It may also comprise at least one memory component, 8905, coupled to the search processor. There are many different system configurations that may be created with the content search processor of this invention. Hence the examples in this patent should not be used as limiting the scope to them, rather they are primarily a means to explain the search processor in a few sample usage scenarios. The search processor of this patent may be used on line cards, network adapters or network interface controllers, storage networking cards, IO cards, motherboards, control processing cards, switching cards or other system elements of systems like networking devices such as routers, switches, management devices, security devices, gateways, virtualization devices, storage networking devices, servers, storage arrays, and the like. The search processor or its components may also be coupled to the microprocessors, network processors, switching chips, protocol processors, TCP/IP processors, control plane processors, chipsets, control processors or

other devices, including being incorporated as a functional block on these processors or chips. The content search processor may be used to perform content inspection at high line rates in the systems in which it is incorporated to offload or assist in content processing to the main processors of such systems. There may be configurations where multiple search
5 processors may also be incorporated in systems to provide scaling in performance or number of rules or a combination thereof for content search. The search processor may be incorporated on network line cards, in line with the traffic and offer line rate deep packet inspection as discussed earlier.

The configuration illustrated in Fig. 89 may be used for email security or instance message
10 security or outbound security or extrusion detection or HIPAA compliance or Sarbanes-Oxley compliance or Gramm-Leach-Bliley compliance or web security or the like or a combination thereof. The security capabilities listed may comprise anti-spam, anti-virus, anti-phishing, anti-spyware, detection/prevention of directory harvest attacks, detection/prevention of worms, intrusion detection/prevention, firewalls, or the like or detection/prevention of leaks of
15 confidential information, health care information, customer information, credit card numbers, social security numbers or the like or a combination thereof. The content search processor or processors in such device may be configured with a set of security rules for one or more of the applications listed above and provide acceleration for content search for information incoming or outgoing from the device. Since, the content search processor may be used
20 inline with the traffic, the device may be deployed at any place in the network, like close to a router or a switch or gateway of an organization's networks or at a departmental level or within a datacenter or a combination and provide high speed content inspection to incoming or outgoing traffic flow of the network.

Fig. 90 illustrates example anti-spam performance bottleneck and solution. As discussed
25 earlier, content search performance using a DFA or NFA based search on a microprocessor results in below 100Mbps performance. Fig. 90 illustrates an anti-spam application as an example application to show the value of hardware based content search. The performance numbers are not illustrated to scale. The figure illustrates four vertical stacks of operations in four types of appliances. The first stack, 9000, is illustrated to represent an email
30 appliance stack. An email appliance typically may comprise device drivers to drive the hardware devices on the appliance, the networking protocol stack along with other functions of the Operating System (OS) and a mail transport agent (MTA) which are all typically software components along with other application software. Today's servers, which are typically used for email appliances, are able to keep up with network line rates of up to
35 1Gbps, and perform the application functions due to the high performance processors.

Typically a 1GHz processor is required to process 1Gbps line rate traffic for network protocol stack processing. Since the state of art processors are around 4GHz today, the servers can handle the network traffic and have processing power available to do other needs of the OS and the applications running on a server. Thus the email appliance stack, 9000, running on
5 a high end server, should be able to keep up with a high line rate. A study by network world magazine, "Spam in the Wild: Sequel" done in December 2004, showed the performance comparison of a large number of anti-spam software and appliance vendors. Under their configuration the range of the message processing performance of the vendor products listed was from around 5 messages per second to 21 messages per second. When this
10 performance number is translated into line rate performance using the worst case message sizes used by network world of 10,000 characters per message, the line rate performance comes to be below 2 Mbps sustained performance. All the vendors either software or appliance solutions were based on dual Xeon processor servers. Thus, a server that can handle 1Gbps network line rate traffic, when performing anti-spam application its
15 performance drops down to below 10 Mbps. The reason for this is that one of the features used extensively by most anti-spam vendors is searching of emails against a set of rules, which are typically represented as regular expressions. The anti-spam appliance stack, 9001, illustrates the email appliance with anti-spam capability loaded on it. Anti-spam applications typically performs many complex regular expression rules based filtering along
20 with statistical filtering, reputation based filtering and the like. The anti-spam rules are typically applied sequentially to each incoming email one after the other to find a rule that may match the content of the email. Then the anti-spam application may apply scores to the rules that match and then decide if a message is spam or not based on the total score it receives. Such an operation causes the stack performance needs to grow substantially
25 higher than a typical email appliance stack, where the anti-spam filters, 9005, overhead on the performance of the appliance is substantial to reduce the over all anti-spam server appliance performance to be below 10Mbps. The content search processor of this invention can be used in such anti-spam appliances to achieve significant performance improvements. The hardware accelerated anti-spam appliance stack, 9002, illustrates the impact of using
30 the search processor of this invention on the overall performance of the system. In such a case, all the anti-spam filters, 9011 thru 9013, may be configured on the search processor, 9006, which in turn may be used to inspect each incoming message. Since all rules would be searched simultaneously, the search processor based appliance can achieve 1Gbps line rate performance or more, since the host CPU is relieved from the performance intensive
35 regular expression searches. The compute device illustrated in Fig. 89 may be one such configuration that may be used as the anti-spam appliance to achieve multiple orders of magnitude higher performance than a standard server based anti-spam appliance. The

stack, 9003, illustrates a stack of an enhanced messaging appliance which may use the TCP/IP offload processor for offloading the protocol processing from the host CPU along with the content search processor of this invention. Thus a significant amount of CPU bandwidth can be made available to other applications which may not have been possible to execute on the computing device without significant performance impact. The use of TCP/IP offload and content search processor may be done individually or in combination and the use of one does not require the use of the other. TCP offload and content search processors could be on the same device providing network connectivity and the acceleration.

Fig. 91 illustrates an anti-spam with anti-virus performance bottleneck. This figure is very similar to Fig. 90, except that the anti-spam appliance whose stack is illustrated also supports anti-virus capability. Anti-virus searches are different than the anti-spam searches but they also add a significant performance overhead as illustrated by the stack, 9104. The number of filters for anti-virus is lot larger than those for anti-spam, though when a content search processor(s) of this invention is used the anti-virus overhead can also be substantially reduced as illustrated by 9105.

Fig. 92 illustrates application content search performance bottleneck and solution. The content search processor of this invention can be used as a search accelerator for a large number of applications that require content search but do the search on the host processor. Since, the performance of the host processors for content search is not very high as discussed above, a content search processor based accelerator can substantially increase the performance of these applications. The applications that require content search are many like data warehousing applications, database applications, bioinformatics related applications, genetics, proteomics, drug discovery related applications and the like. The figure illustrates three boxes, 9200, 9201 and 9202 which represent the content search based application performance in terms of host CPU load. The traditional applications run on a server or a workstation or personal computer, and perform content search interspersed with other tasks that the application needs to perform. If the applications perform a significant amount of search, then the performance need of the search portions of the application can be substantially higher than the other parts. This is illustrated by content search portions of applications app1 and appN, 9203 and 9205 respectively, compared to the other code of these applications, 9204 and 9206. The stack in 9200 is the how the current or prior art solution exists for content search applications. Though the stack illustrates a continuous stack for content search and other code sections, the actual execution may generally be composed of search interspersed with other code functions. When a content search processor and accelerator of this invention is used in the computing

device performing this function, it may be possible to have the application leverage the search capabilities of the processor and accelerate the performance of the application substantially compared to a computing device without the search acceleration support. The stack in 9201, illustrates the impact on the CPU load and the resulting time spent by the application when converted to leverage the content search accelerator. The stacks 9203 and 9205, could take substantially less load and time as illustrated by stacks, 9207 and 9208 respectively. Similarly, the performance of the system may be further increased by offloading the TCP/IP protocol processing as illustrated by 9209. As described above, TCP/IP offload and content search offload are independent of each other and may each be done without the other in a system. However, one could also use the content search processor with the TCP/IP processor together as separate components or on the same device and achieve the performance benefits.

Fig. 93 illustrates an example content search API usage model. As discussed above, the content search processor may be used to accelerate content search portions of generic applications. To enable an ease of creation of new applications and migrate existing applications to leverage such search processor acceleration capability one may create an application programming interface (API) for content search. An example content search API is illustrated in Fig. 94 and described below. The content search API may reside in the user level or the kernel level with user level calls, or a combination. The Fig. 93 illustrates the content search API at the user layer, 9307. The content search API would provide API functions that any application can call to get the benefit of content search acceleration. There would be a convention of usage for the applications to use the content search API. For example the application may be required to setup the search rules that can be configured on the search processor using the API calls before the application is run or may be required to dynamically create the rules and setup them up in the appropriate format so that they can be configured on the content search processor using the API or a combination. There would be API calling conventions that may be established dependent on the hardware system, the operating system or the search processor or a combination. The applications may then be coded to the API conventions and benefit from the search processor acceleration. The figure illustrates applications App1, 9300 through App N, 9303, working with the content search API, 9307 to get access to the content search acceleration hardware, 9317, using logical interface paths illustrated as 9312, 9313 and 9314. The content search API may access the services and resources provided by the content search accelerator through a port driver which may be running under a kernel. The applications may pass the content to be searched directly through this interface or put the content to be searched as well as tables to be setup as needed, in the application's buffers, 9304, 9305,

and then instruct the content search processor to retrieve the information from these buffers through the content search API. The API may map these buffers to the kernel space so the port driver for the search API can provide them to the content search processor or the buffers may be made available for direct memory access by the search processor hardware.

- 5 The search processor may store the content in on-chip or off-chip memory buffers, 9318, and then perform the requested search on the content. Once the search is complete the results of the search may be provided back to the application using a doorbell mechanism or a callback mechanism or data buffers or the like as allowed by the operating systems' model. The content search API may provide a polling mechanism as well which may be used by the application to check and/or retrieve the search results.
- 10

Fig. 94 illustrates an example content search API with example functions. The figure illustrates a set of functions which may be a part of the example content search API.

- Though, the list of functions may be more or less than those illustrated, the functions provide a basic set that would enable an application to use the content search hardware with the use of the API. The example functions do not illustrate the input, output or return parameters for API function calls, which may depend on the operating system, calling conventions and the like. An application may use the API, by first querying the capabilities of the hardware engine and then initializing it with appropriate rules, pointers, permissions and the like that may be required for the content search processor to communicate with the application and its resources through the kernel or the user mode or a combination. The application may set specific rules as DFA rules or NFA rules which may get configured in the search processor. An application may be given access to multiple contexts that it may be able to leverage to perform context based search. The application can start performing the search against its content once the content search processor is appropriately setup with all necessary rules.
- 15
- 20 The application can communicate the content to be searched directly to the search processor using the API by sending byte stream through the interface. There may be versions of an API function, not illustrated, like sendData() which may be used by an application to start sending data to the search processor, start the search and to indicate when the search processor should stop searching. A more efficient way of performing the search may be that the application may fill a buffer or a set of buffers to be searched, and then provide the search processor with a pointer(s) to the buffer(s) who can then start searching the buffers with the configured rules once it receives a call to start the search using an API call like startHWsearch(). The search processor may have been initialized to communicate the results of the search to the application through one of many mechanisms like copying the results to a result buffer or storing the result on the memory associated with the search processor or invoking a callback function registered by the application to the
- 25
- 30
- 35

operating system or the like. The search processor may also communicate to the application with a doorbell mechanism to inform it that the search is done. There are many different ways of communicating the information as described earlier and may be dependent on the operating system and the system hardware architecture. There may also be polling
5 mechanism available with an API function like `isSearchDone()`, not illustrated, which may provide the answer to a query to the search hardware whether a specific search is complete. If the answer from the hardware to the application is that the search is done, then the application may ask for the specific result using an API call like `getRes()`, or the application may ask for a pointer to a buffer that may hold the result using a call like `getResPtr()`
10 illustrated in Fig. 94. Once the application is done with the specific search or is done using the search processor it may call the API function `stopHWsearch()` to stop the hardware processor from performing the search for this application. There may also be an API call like `removeAppContext()`, not illustrated, which may be called by the application to indicate to the OS and the search processor hardware that the application is done using the search
15 processor and hence all its associated context may be freed-up by the search processor hardware for use by another application that may need the search processor resources. There may be other hardware features specific API calls as well, like `setRuleGroup()`, `selectRuleGroup()`, `setInitGroup()` and the like, that may allow an application to create groups of rules and the order of their execution using mechanisms of rule grouping
20 described earlier. As discussed earlier there may be many more functions and variation of API functions that can be created to enable a general content search application acceleration using a hardware search processor from the teachings of this patent that will be obvious to one skilled in the art. Thus it is possible to create a content search API to provide content search capabilities to general applications. Though, the description above is given
25 with an example where the rules to be used are setup by an application before starting the search, it may be possible to update the rule set that is configured in the search processor dynamically while the search is in progress by adding, removing and/or modifying the rules that have already been configured to start using the updated rule set for any future searches by the application.

30 Fig. 95 illustrates an example application flow (static setup) using the search processor. The flow illustrates a static process for setting up the rules and the hardware processor although as discussed above a dynamic setup is also feasible as would be obvious to one skilled in the art. The flow may allow an application to add/remove/modify rules in the processor as the application executes at runtime to enable a dynamic flow. The illustration provides a
35 mechanism where existing applications or new applications may be updated with content search rules and API calls which can enable the application to use a content search

processor. An application source, 9500, may be updated, 9501 to create application source with modifications for content search where the content search rules may be setup in distinct code sections or may be clearly marked, 9502, as expected by the compiler coding conventions, which is then compiled by a content search aware compiler, 9503. The
5 compiler generates an object code, 9504, with content search rules compiled in sections which a loader may use to configure them in the search processor. The application object code may then be distributed to customers or users of content search processor for accelerating the application's search performance. The application code may be distributed electronically using the internet, world wide web, enterprise network, or other network or
10 using other means like a CD, DVD, or other computer storage that can be used to load the application. The application update, 9501, may be done manually or using a tool or both as appropriate. The distributed object code, 9506, is read by the loader, 9507, or a similar function provided by an operating system to which the application is targeted, and setup for execution on the system. The loader or another function may use a set of API calls or a port
15 driver or other OS function or a combination to configure the content search processor with appropriate rules that the application needs as coded in the object code as illustrated by block 9508. Once the search application hardware is setup and other resources that the application needs get reserved or setup, the application is started, 9509, by the OS. The application may execute or perform tasks, 9510, if needed before content search. The
20 application may then setup the content, 9511, it needs to search by the search processor. Then it starts the search processor to perform search, 9513. Once the search is done it may retrieve the results, 9514. While the search is being conducted by the search processor, the application may continue to perform other tasks, 9512, on the main CPU or other elements of the system. If the application is done the application may exit, otherwise the application
25 may continue the execution where more tasks may be performed including new search if necessary. The flow diagram illustrates the execution of tasks as a loop from 9515 to 9510, though the tasks being executed may be very different from one time to the next through the loop. The loop is not illustrated to mean that the same code sequence is being repeated. It is meant to show that the type of tasks may be repeated. Further, not all tasks from 9510
30 through 9515 may need to be present in an application flow. Once the application is done, it may release all the resources it uses beside those for the content search processor.

Fig. 96 illustrates FSA compiler flow. The content search rules in a regular expression form are used as an input to the FSA compiler flow. The FSA compiler searches through all the rules or groups of rules, and parses them, 9602. The compiler then generates an NFA and a
35 DFA to meet the content search processor's implemented architecture as illustrated by steps 9603 and 9604. The compiler then compares the cost of implementing the RE as an NFA

- and a DFA and selects the NFA or DFA according to the processor capabilities as illustrated by the steps 9606 and 9611. Compiler generates the NFA or DFA code possibly with action and/or tag as needed by the RE or RE group. The compiler continues from the steps 9602 through 9607 until all the RE rules have been processed. It then generates all the rule FSA code and tags as necessary for the entire RE rule set as illustrated by 9608 and stores the rules in a rules database in a storage, 9609, which may be RAM, ROM, hard disk, floppy, CD, DVD or any other computer storage. The compiled rules may then be distributed to the content search processors as illustrated by the step 9610. The rules may be distributed by a distribution engine like that described above for the security compiler flow.
- Fig. 97 illustrates a search compiler flow (full and incremental rule distribution). The flow can be used for distributing search rules or security rules when the full set of rules are defined or when any updates or modifications are made to the rule set and incremental changes to the rule set need to be communicated and configured in the search processor. The flow illustrated in Fig. 97 is very similar to that illustrated in Fig. 67 except that the FSA compiler flow described above is a part of the flow in Fig. 97. The rules like application layer rules, network layer rules or storage network layer rules or any other search rules may be created using manual or automated means and provided as inputs to the search compiler flow in a predefined format. The search compiler's rule parser, 9704, parses the rules and converts them into regular expression format if the rules are not already in that form. Then the regular expression rules are passed through the FSA compiler flow which is illustrated in Fig. 96. The output of the FSA compiler is the rules compiled to the node capabilities of the node that has the content search processor and stored in the rules database. The rules from the rule database are retrieved and distributed by the rules distribution engine to the appropriate node(s) with the search processor. The rules distribution engine, block 9709, distributes the rules to the appropriate nodes using a central manager flow and rules distribution flow similar to that illustrated in Fig. 57 and Fig. 58. The search or security rules may be distributed to the host processor or a control plane processor as illustrated in Fig 58 or to a control processor and scheduler, block 8402, described above, or a combination thereof as appropriate depending on the node capability. The rules may be distributed using a secure link or insecure link using proprietary or standard protocols as appropriate per the specific node's capability over a network. The network may be a local area network (LAN), wide area network (WAN), internet, metro area network (MAN), wireless LAN, storage area network (SAN) or a system area network or another network type deployed or a combination thereof. The network may be Ethernet based, internet protocol based or SONET based or other protocol based or a combination thereof.

Fig. 98 illustrates FSA synthesis and compiler flow for FPGA. The Flow illustrated in this figure is similar to the FSA compiler flow illustrated in Fig. 96. The steps, 9801 through 9811 are essentially the same as the steps 9601 through 9611. The FSA synthesis and compiler flow for the FPGA illustrates a few additional steps which can be used to convert the FSA rules to synthesize a content search processor for them in an FPGA. The RE rules converted to NFA or DFA is used to generate appropriate RTL in an automated way as illustrated by 9812. The NFA or DFA RTL generation may use NFA and DFA building block RTL from a library as illustrated by 9817. Once the RTL for all elements of the search processor array is generated from all the RE rules, the top level RTL for the search processor may be generated with NFA or DFA specific search engines targeted for the compiled rules. The output RTL is then merged with any other RTL code for the search processor, for example a memory controller, and output may be stored in computer storage. This merged RTL is then provided as an input to the FPGA RTL synthesis tools which generate an FPGA based search processor specifically targeted to meet the needs of the regular expression rules that are input to the FSA synthesis and compiler flow.

The processors of this invention may be manufactured into hardware products in the chosen embodiment of various possible embodiments using a manufacturing process, without limitation, broadly outlined below. The processor may be designed and verified at various levels of chip design abstractions like RTL level, circuit/schematic/gate level, layout level etc. for functionality, timing and other design and manufacturability constraints for specific target manufacturing process technology. The processor design at the appropriate physical/layout level may be used to create mask sets to be used for manufacturing the chip in the target process technology. The mask sets are then used to build the processor chip through the steps used for the selected process technology. The processor chip then may go through testing/packaging process as appropriate to assure the quality of the manufactured processor product.

While the foregoing has been with reference to particular embodiments of the invention, it will be appreciated by those skilled in the art that changes in these embodiments may be made without departing from the principles and spirit of the invention.

30

Claims:

- 1 1. A runtime adaptable search processor operating on information content and having a
2 plurality of hardware configurations, comprising:
 - 3 a. a control processor and scheduler to control and schedule said information
4 content for processing;
 - 5 b. an adaptation controller to configure said runtime adaptable search processor
6 configurations at runtime;
 - 7 c. at least one configuration memory to store at least said plurality of hardware
8 configurations;
 - 9 d. at least one runtime adaptable NFA search engine for performing search on
10 said information content using at least one NFA;
 - 11 e. at least one runtime adaptable DFA search engine for performing search on
12 said information content using at least one DFA ; and
 - 13 f. a combination of any of the foregoing.
- 1 2. The runtime adaptable search processor of claim 1 further comprising at least one of
2 the following:
 - 3 a. at least one search cluster comprising at least one search engine for
4 searching information;
 - 5 b. a global configuration controller for global configuration control;
 - 6 c. a configuration memory controller for controlling accesses to at least said
7 configuration memory;
 - 8 d. an inter cluster routing switch for controlling routing between a plurality of said
9 at least one search cluster;
 - 10 e. a global memory for providing memory storage;
 - 11 f. a global memory controller for controlling accesses to at least said global
12 memory; and
 - 13 g. a combination of any of the foregoing.
- 1 3. The runtime adaptable search processor of claim 2 wherein said at least one search
2 cluster comprises:

- 3 a. a cluster configuration controller for configuring said at least one search
4 cluster;
- 5 b. at least one runtime adaptable search engine for searching information; and
- 6 c. a configurable cluster interconnect network for connecting a plurality of said at
7 least one runtime adaptable search engine.
- 1 4. The runtime adaptable search processor of claim 2 wherein said at least one search
2 cluster further comprises at least one of the following:
- 3 a. a cluster memory for providing memory storage;
- 4 b. a cluster memory controller for controlling at least said cluster memory;
- 5 c. a cluster routing switch for interconnecting said at least one runtime
6 adaptable search engine to said at least one search cluster or said inter cluster routing
7 switch;
- 8 d. a cluster flow controller for at least storing and retrieving flow context; and
- 9 e. a combination of any of the foregoing.
- 1 5. The runtime adaptable search processor of claim 3 wherein said at least one runtime
2 adaptable search engine comprises at least one of the following:
- 3 a. a runtime adaptable NFA search engine for performing search using at least
4 one NFA;
- 5 b. a runtime adaptable DFA search engine for performing search using at least
6 one DFA; or
- 7 c. a combination of the foregoing.
- 1 6. The runtime adaptable search processor of claim 1 wherein said at least one runtime
2 adaptable NFA search engine evaluates at least one said NFA and comprises:
- 3 a. at least one state logic block to perform state logic operations for a state;
- 4 b. at least one application state memory to store application state information;
- 5 c. at least one symbol detection logic to detect a symbol;
- 6 d. at least one application state controller to control the operation of said at least
7 one runtime adaptable NFA search engine; and
- 8 e. a combination of any of the foregoing.

1 7. The runtime adaptable search processor of claim 1 wherein said at least one runtime
2 adaptable NFA search engine further comprises at least one of:

- 3 a. pattern substitution control to perform pattern substitution;
- 4 b. counters to provide counters for various events;
- 5 c. match logic to detect a rule match;
- 6 d. FSA extension logic to enable said runtime adaptable NFA search engine to
7 support larger FSAs and
- 8 e. a combination of any of the foregoing.

1 8. The runtime adaptable search processor of claim 1 wherein said at least one runtime
2 adaptable DFA search engine comprises:

- 3 a. at least one DFA operations engine to perform state logic operations;
- 4 b. at least one DFA instruction and state flow memory to store DFA operation
5 instruction and state flow of instructions;
- 6 c. at least one DFA state controller to control the operation of said at least one
7 runtime adaptable DFA search engine; and
- 8 d. a combination of any of the foregoing.

1 9. The runtime adaptable search processor of claim 1 wherein said at least one runtime
2 adaptable DFA search engine further comprises at least one of:

- 3 a. pattern substitution control to perform pattern substitution;
- 4 b. counters to provide counters for various events;
- 5 c. a DFA context controller to control the application context of said DFA search
6 engine;
- 7 d. DFA extension logic to enable said runtime adaptable DFA search engine to
8 support larger FSAs; and
- 9 e. a combination of any of the foregoing.

1 10. The runtime adaptable search processor of claim 6 wherein said at least one state
2 logic block comprises:

- 3 a. a state dependent vector (SDV) used to enable or disable interconnect
4 between said state and one or more states of said NFA;

- 5 b. a current state vector (CSV) representing the current state of some of said
6 one or more states of said NFA that said state depends on for its next state;
- 7 c. a received symbol vector (RSV) representing the value of said symbol
8 received as input to said symbol detection logic;
- 9 d. a left-biased or right-biased signal that indicates the NFA type;
- 10 e. a state transition logic that uses at least said SDV, said CSV, said RSV, and
11 said left-biased or right-biased signal to generate the next state value for said state;
- 12 f. an init logic block to control the value of the next state for said state;
- 13 g. a state memory to store the value of said state;
- 14 h. an accept detect logic to detect and signal when said state is an accept state
15 and is activated; and
- 16 i. a combination of any of the foregoing.

1 11. The runtime adaptable search processor of claim 6 wherein said at least one state
2 logic block further comprises at least one of the following:

- 3 a. a programmable storage for said SDV to enable runtime configuration of said
4 SDV;
- 5 b. a programmable storage for said left-biased (LB) or right-biased (RB) signal
6 to enable runtime configuration of a realized NFA type;
- 7 c. a programmable storage for storing a start flag to indicate if said state is a
8 start state of said NFA;
- 9 d. a programmable storage for storing an accept flag to indicate if said state is
10 an accept state of said NFA;
- 11 e. a programmable storage for storing an action field to indicate the action that
12 should be taken when said state is activated or reached during said NFA evaluation;
- 13 f. a programmable storage for storing a tag field to indicate a tag that is issued
14 when said state is activated or reached during said NFA evaluation; and
- 15 g. a combination of any of the foregoing.

1 12. A runtime adaptable finite state automaton architecture having a plurality of states,
2 comprising:

- 3 a. at least one interconnect between any two of said plurality of states;

- 4 b. a state dependent vector (SDV) for each of said plurality of states that is used
5 to enable or disable of said at least one interconnect between said two states and other
6 states of said plurality of states;
- 7 c. a current state vector (CSV) for each of said plurality of states representing
8 the current state of some of said plurality of states which the state of said plurality of states
9 depends on for its next state evaluation;
- 10 d. at least one symbol associated with each of said plurality of state to control
11 the transition into or out of said each of said plurality of states;
- 12 e. at least one symbol detection logic for receiving symbols, detecting the value
13 of said symbols and generating a received symbol vector (RSV);
- 14 f. at least one application state memory to hold an application context that uses
15 said finite state automaton;
- 16 g. a state transition logic for each of said plurality of states that uses at least
17 said state dependent vector, said current state vector, and said received symbol vector to
18 generate the next state for each of said plurality of states; and
- 19 h. a combination of any of the foregoing.

1 13. The runtime adaptable finite state automaton (FSA) of claim 12 further comprises at
2 least one of the following:

- 3 a. a programmable storage for SDV to enable runtime configuration of SDV;
- 4 b. a programmable storage for a left-biased (LB) or right-biased (RB) signal to
5 enable runtime configuration of the realized FSA type;
- 6 c. a programmable storage for storing a start flag to indicate if said state is a
7 start state of said FSA;
- 8 d. a programmable storage for storing an accept flag to indicate if said state is
9 an accept state of said FSA;
- 10 e. a programmable storage for storing an action field to indicate the action that
11 should be taken when said state is activated or reached during said FSA evaluation;
- 12 f. a programmable storage for storing a tag field to indicate a tag that is issued
13 when said state is activated or reached during said FSA evaluation; and
- 14 g. a combination of any of the foregoing.

1 14. A system for operating on network packets or on local content, said system capable
2 of being coupled to a network, said system comprising a runtime adaptable search
3 processor to process content embedded within said network packets or within said local
4 content.

1 15. The system of claim 14 wherein said runtime adaptable search processor further
2 provides runtime adaptable search processor array hardware to dynamically select hardware
3 configurations within said runtime adaptable search processor to support application or
4 service processing needs of said network packets or said local content.

1 16. The system of claim 14 comprising a content search API to enable applications to
2 utilize the runtime adaptable search processor hardware to accelerate content search;
3 wherein the content search API comprises means to initialize said runtime adaptable search
4 processor; means to communicate content search rules to said runtime adaptable search
5 processor; means to communicate the content to be searched by said runtime adaptable
6 search processor; means to communicate the results of said content search by said runtime
7 adaptable search processor to said applications; means to start and stop said runtime
8 adaptable search processor to start and stop content search; and a combination of any of
9 the foregoing.

1 17. The system of claim 14 wherein said runtime adaptable search processor enables
2 email security applications comprising anti-spam; anti-virus; anti-spyware; anti-phishing;
3 confidential leak prevention; or regulatory compliance; or a combination of the foregoing; or
4 web security applications comprising firewall; intrusion detection and prevention; application
5 layer security; anti-spam; anti-virus; anti-spyware; anti-phishing; confidential leak prevention;
6 or regulatory compliance; or any combination of the foregoing; wherein said applications can
7 be performed at network line rates; wherein said network line rates lie within a range of line
8 rates from below 100Mbps to 10Gbps and higher.

1 18. The runtime adaptable search processor of claim 6 wherein said at least one
2 application state memory comprises at least one application state context that represents the
3 context information to be programmed in said NFA search engine when said at least one
4 application state context is selected to be configured in said NFA search engine; wherein the
5 application state context comprises NFA parameters which are applicable to all states of
6 said NFA, and at least one set of state parameters being applicable to each state of said at
7 least one state of said NFA or a combination of the foregoing.

1 19. The system of claim 14 wherein said runtime adaptable search processor receives
2 data packets from a network and sends data packets to a network, and further provides
3 capability to select:

- 4 a. policy based services, applications or policies or a combination thereof;
- 5 b. user selected services, applications or policies or a combination thereof;
- 6 c. user defined services, applications or policies or a combination thereof;
- 7 d. user configured services, applications or policies or a combination thereof;
- 8 and
- 9 e. a combination of any of the foregoing;

10 on the data packets received from said network, on the data packets sent to said network, or
11 on a combination of thereof.

1 20. The runtime adaptable search processor of claim 1 embedded in a system for
2 operating on network packets or on local content, said system capable of being coupled to a
3 network, wherein said runtime adaptable search processor processes content embedded
4 within said network packets or within said local content, said runtime adaptable search
5 processor further comprising runtime adaptable search processor array hardware to
6 dynamically select hardware configurations within said runtime adaptable search processor
7 to support application or service processing needs of said network packets or said local
8 content.

Fig. 1 SCSI Architecture Layers




Fig. 2 SCSI Architecture Layers on iSCSI




Fig. 3 OSI Stack




Fig. 4 OSI Stack (w/ HW TCP/IP)




Fig. 5 Host Software Stack (iSCSI)




Fig. 6 S/W TCP Stack Data Transfer




Fig. 7 Remote Direct Memory Access




Fig. 8 Host Software Stack (SCSI)




Fig. 9 iSCSI Stack




Fig. 10 TCP/IP Stack




Fig. 11 iSCSI Data Flow




Fig. 12 iSCSI Data Structures




Fig. 13 TCP/IP Session DB Entry

Source_IP	Destination_IP	Source_Port	Destination_Port	Protocol
Connection_ID	TCP_Window_PTR	Window_Size	Sender MSS/MTU	Expected_SEQN
TCP_State	TCP_SND_SEQ#	TCP_SEQ#	TCP_ACK#	TCP_RCV_SEQ#
Fragment_PTR	Packet_send_PTR (byte granular)	Packet_Start_Time	Last_PKT_Time	TCP_Parameters (slow start, retransmit..)
Packets_Count	Congestion Window	SSThresh	Receiver MSS	Other Fields
Frame_key	Out of Order PTR	Other Frame Fields		
Hash_Key	NXT_H_Match_PTR	Valid/Invalid		

Fig. 14 iSCSI Session DB Entry

iSCSI_State	Data_SEQ#	Command_SEQ#	Task_TAG	R2T_SEQ#
Status_SEQ#	EXP_Status_SEQ#	EXP_CMD_SEQ#	TARGET_MODE (Solicited or not)	Target_Parameters first burst, max PDU ...)
ISID	TSID	Portal Group Tag	Transfer Direction	Other fields
Connection_ID	Command ID	RDMA enabled	RDMA Keys	Other RDMA Param
MDL_List_Pointer	Mem Descriptor	Block size	No. of Blocks (xmitted or rev'd)	ULP_Pointers
Hash_Key	NXT_H_Match_PTR	Valid/Invalid		

Fig. 15 iSCSI Session Memory





Fig. 16 IP Network Application Processor






Fig. 17 IP Network
Application Processor

Fig. 18 Input Queue & Controller




Fig. 19 Packet Scheduler & Sequencer




Fig. 20 Packet Classification/Policy Engine




Fig. 21 SAN Packet Processor




Fig. 22 SAN Packet Processor




Fig. 23 TCP/IP Processor




Fig. 24 IP Storage Engine (iSCSI Processor)




Fig. 25 Output Queue




Fig. 26 Storage Flow & RDMA Controller




Fig. 27 Host Interface Controller





Fig. 28 Security Engine

Fig. 29 IP Storage Session Cache &
Memory Controller Complex




Fig. 30 Classification Engine Data Structure




Fig. 31 Read Operation




Fig. 32 Read Data Packet Flow




Fig. 33 Write Operation




Fig. 34 Write Data Packet Flow




Fig. 35 iSCSI Read using RDMA Write




Fig. 36 Read Data Packet Flow (RDMA)




Fig. 37 iSCSI Write using RDMA Read






Fig. 38 Write Data Packet Flow (RDMA)









Fig. 40 Read Data Flow






Fig. 42 Read Data Flow (Secure)




Fig. 43 R2T &
Write Data Flow (secure)




Fig. 44 RDMA
Buffer Adv




Fig. 45 RDMA Write Flow





Fig. 46 RDMA
Read Flow







Fig. 48 Session
Tear down






Fig. 50 Target Write
Flow (Send R2T)




Fig. 51 Target Write Flow (Receive Write)




Fig. 52 Target
Read Flow




Fig. 53 Enterprise Network with Perimeter Security





Fig. 54 Enterprise Network with Distributed Security

Fig.55 Enterprise Network with Distributed Security






Fig. 56 Central Manager/Policy Server & Monitoring Station




Fig. 57 Central Manager Flow




Fig. 58 **Rule Distribution Flow**




Fig. 59 Control Plane Processor/Policy Driver Flow

Fig. 60 Packet Filtering Rules Example

Direction	Source	Dest	Prot	Src. Port	Dst Port	Other (e.g. App Layer, User name, UPL Prot etc.)	Action
In	Ext	Int	TCP	> 1023	25	Any	Permit
Out	Int	Ext	TCP	25	> 1023	Any	Permit
In	Int	Int	TCP	Any	80	From 8am to 5pm	Permit
Out	Int	Int	TCP	80	Any	Any	Permit
In	Ext	Int	ICMP	Any	Any	Any	Deny
In	Ext	Int	TCP	Any	80	Long URL in payload	Deny
Out	Int	Int	UDP	5004	>1023	Any	Permit
Out	Int	Ext	UDP	5004	>1023	Any	Deny
Out	Int	Ext	UDP	80	Any	Social Sec Numbers	Deny
Either	Any	Any	Any	Any	Any	Any	Deny





Fig. 61 TCP/IP Processor






Fig. 63 Adaptable TCP/IP Processor Alternate




Fig. 64 Runtime Adaptable Processor




Fig. 65 Compute Cluster




Fig. 66 Security solution




Fig. 67 Security solution compiler flow




Fig. 68 Flow-through Secure Network Card




Fig. 69 Network line card with look-aside security




Fig. 70 Security and content search acceleration adapter




Fig. 71 Security processor




Fig. 72 An alternate Security processor




Fig. 73 A third Security processor




Fig. 74 Content Search and Rule Processing Engine

Rule #	Rule Group	Regular Exp	Action/Tag
R1:	<COMMENT>	[^\\-\\\"\\n\\r]+	/* Skip the comment */
R2:	<COMMENT>	-\\>	/* Skip the comment */
R3:	<COMMENT>	-{2,}\\>	START(INITIAL);
R4:		[^\\\"\\>\\r\\n]+	/* Skip the FONT tag */
R5:		\\>	START(INITIAL);
R6:	<S_TAG>	!-	START(COMMENT);
R7:	<S_TAG>	FONT	START(FNT);
R8:	<S_TAG>	\\>	START(INITIAL);
R9:	<INITIAL>	\\<	START(S_TAG);

Fig. 75 An example of Regular Expression Rules (Prior Art)

Regular Expression: $(xy + y)*yx$




Fig. 76a Thompson's NFA (Prior Art)




Fig. 76b Berry-Sethi NFA (Prior Art)






Fig. 76c DFA (Prior Art)






Fig. 78a State Transition Controls

NS = A

PS	Sym	SDV
A	S1	V1
B	S1	V2
C	S1	V3
D	S1	V4

7811

Left Biased

NS = A

PS	Sym	SDV
A	S1	V1
B	S2	V2
C	S3	V3
D	S4	V4

7812

Right Biased

State Prop: LB, RB, SDV, Start, Accept, Action
Fig. 78b Configurable Next State Table per State




Fig. 79a State transition logic (STL) for a state




Fig. 79b A State Logic Block




Fig. 80 A NFA based Search Engine




Fig. 81 Application State Memory Example Configuration




Fig. 82 A DFA based Search Engine

Operations	RE Equiv	Description
EQ X	X	Go to next state if Input is X otherwise go to alternate state
NEQ X	~X	Go to next state if Input is not X otherwise go to alternate state
RANGE X, Y	[X-Y]	Go to next state if Input is in Range X to Y else alternate state
LU_EQ X	. *X	Wait in the current state until first occurrence of X then go to next
LU_NEQ X	X+	Wait until the last occurrence of X in a series then go to next state
LU_MIN_EQ X	[n..] X	Wait in current state and Detect Min 'n' occurrences of X. If less than 'n' occurrence go to alternate state
LU_MIN_MAX_EQ X	[n..m] X	Wait in current state and Detect 'n' to 'm' occurrences of X. If not in the range go to alternate state
EQ_CHAR	[a-zA-Z]	Go to next state if Input is in alphabet range else alternate state
EQ_DIGIT	[0-9]	Go to next state if Input is a decimal digit else alternate state
GEN_FLAG		Generate action or tag and a flag that the accept state is reached

Fig. 83 Example DFA Operations




Fig. 84 A Runtime adaptable search processor







Fig. 85 A Runtime Adaptable Search Processor Array 8512











Fig. 89 Computing Device with Content Search Accelerator







Fig. 91 Anti-spam with Anti-virus Performance* Bottleneck






Fig. 93 An example content search API usage model

API *	Description
(*parameters are not illustrated)	
Initialize()	Initialize the hardware
getCapabilities()	Query the capabilities of the current hardware
setDFARule()	set a DFA rule in the hardware processor
setNFARule()	set an NFA rule in the hardware processor
setContext()	set the context for the search engine(s)
sndDataPtr()	send the pointer of the data buffer that needs to be searched
getRes()	retrieve the result from the hardware processor
getResPtr()	retrieve the pointer to the memory buffer that holds the results
switchContext()	Switch the current active context of the DFA/NFA engines
setRuleGroup()	set the rule group that a specific rule belongs
selectRuleGroup()	select a specific rule group for access
setInitGroup()	set the initial group that is active for a given flow
startHWSearch()	start the hardware search engine
stopHWSearch()	stop the hardware search engine

Fig. 94 An example Content Search API with example functions





Fig. 95 An example application flow (static setup) using search processor

**Fig. 96 FSA Compiler Flow**




Fig. 97 Search compiler flow (full + incremental rule distribution)




Fig. 98 FSA Synthesis & Compiler Flow for FPGA