

(19) World Intellectual Property
Organization
International Bureau



(43) International Publication Date
9 September 2005 (09.09.2005)

PCT

(10) International Publication Number
WO 2005/081855 A2

(51) International Patent Classification: **Not classified**

(21) International Application Number:
PCT/US2005/005159

(22) International Filing Date: 18 February 2005 (18.02.2005)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
10/783,890 20 February 2004 (20.02.2004) 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 Rudnick Gray Cary 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,

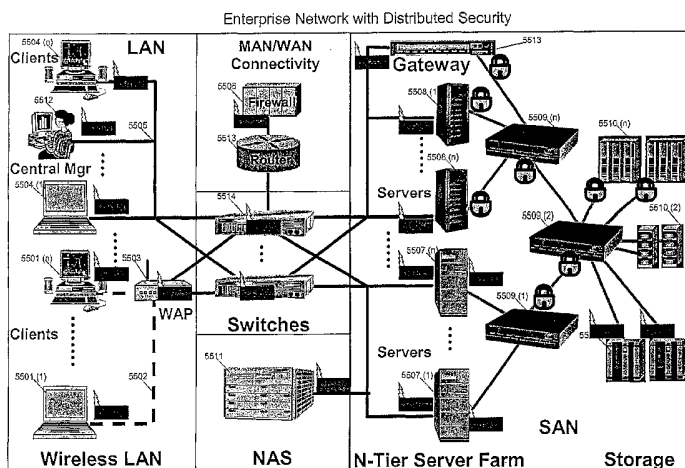
AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SM, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, 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, 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: A DISTRIBUTED NETWORK SECURITY SYSTEM AND HARDWARE PROCESSOR THEREFOR



(57) Abstract: An architecture 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. 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 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 2005/081855 A2

A DISTRIBUTED NETWORK SECURITY SYSTEM AND A HARDWARE PROCESSOR THEREFOR

RELATED APPLICATIONS

Priority is claimed to U.S. Application Serial No. 10/783,890, filed on February 20, 2004.

5 BACKGROUND OF THE INVENTION

This invention relates generally to storage networking semiconductors and in particular to a high performance network storage processor that is used to create Internet Protocol (IP) based storage networks.

10 Internet protocol (IP) is the most prevalent networking protocol deployed across various 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
15 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
20 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.

25 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
30 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 will create a need for enabling security processing in

hardware inside core or end systems beside a perimeter firewall as one of the prominent means of security.

SUMMARY OF THE INVENTION

I 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 processing including firewall, encryption, decryption, intrusion detection and the like.

10 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 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 the host processor.

20 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 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 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 1 Gbps, the memory size overhead and memory speed bottleneck resulting from these add 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 packets will cause severe performance degradation due to context switching overhead,
5 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
10 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
15 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
20 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
25 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.

30 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
35 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.

5 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.

10 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
15 creates potential threats to the information assets inside an enterprise network. Recent research by leading firms and FBI found that over 70 percent of the unauthorized access to information systems is committed by employees or trusted insiders and so are over 95 percent 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
20 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
25 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
30 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

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 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.

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 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. 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 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 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
5 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.

10 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
15 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
20 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
25 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
30 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, LUN masking and the like.

BRIEF DESCRIPTION OF THE DRAWINGS

- 5 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 -
10 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.
- 15 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.
- 20 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.
- 25 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 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.

5 Fig. 16 illustrates a high-level architectural block diagram of an IP network application 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.

10 Fig. 19 illustrates a packet scheduler, sequencer and load balancer useful in one 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.

15 Fig. 22 illustrates an embodiment of the SAN packet processor block of the described IP 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.

20 Fig. 24 illustrates an embodiment of the programmable IP Storage processor engine which 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.

25 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.

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.

5 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.

10 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.

15 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.

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

20 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.

25 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 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.

5 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 processor.

10 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.

15 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.

20 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.

25 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.

Fig. 60 illustrates a sample of packet filtering rules that may be deployed in distributed security systems.

DESCRIPTION

I provide a new high performance and low latency way of implementing a TCP/IP stack in
5 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
10 enterprise networks in a distributed manner using a hardware TCP/IP implementation with appropriate security capabilities in hardware having processing elements to enable high performance and low latency IP based network security applications. The hardware processor 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,
15 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 processor may also be used inside an I/O chipset of one of the end systems.

This system enables distributed security capabilities like firewall, intrusion detection, virus
20 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
25 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 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.

30 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, security capabilities, programmable classification and policy processing features 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 providing hardware based TCP/IP functions or hardware TCP/IP and security functions that relieve the host server processor or processors of TCP/IP and/or security software and performance overhead. The IP processor may be a companion processor to a server chipset, providing the high performance networking interface with hardware TCP/IP and/or security. Servers can be in various form factors like blade servers, appliance servers, file servers, thin servers, clustered servers, database 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 motherboard.

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.

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.

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.

Processor Chipset Component

5 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
10 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.

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
15 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 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
20 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 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
25 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 enterprise storage area networks as well.

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.

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.

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.

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
5 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
10 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
15 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
20 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
25 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
30 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.

Detailed Description

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 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 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 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 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 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 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 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 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.

10 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 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
15 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 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
20 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 network that eliminates or greatly reduces the performance penalty on the host processor, and the resulting latency impact.

25 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 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
30 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 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
35 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 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 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 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 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 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 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 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.

5 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
10 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.
15 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
20 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.

25 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,
30 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
35 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 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 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 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. 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 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 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, 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 T, 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.

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 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 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 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, 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 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. 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.

[start here]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.

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 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 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.

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 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 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 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

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
5 the last fragment arrives, the fragment controller of block 1806, working with the 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
10 of this processor. The input queue controller of Fig. 18 assigns a packet tag/descriptor to 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
15 through the processor by the scheduler and other elements of the processor in an efficient 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
20 headers from the header queue, block 1902, from the input queue controller 1901 to pass them to the classification engine 1703 of Fig. 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
25 which it routes to the scheduler after classification. The scheduler comprises a classification 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
30 engine of Fig. 17. The scheduler state control and the sequencer, block 1916, receive state 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
35 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

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 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 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 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 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 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 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 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.

5 The classifier can be a programmable processor that examines various fields of 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
10 be a programmable processor as well. The primary differences are the 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
15 appropriate packet fields based on the global field descriptor sets, 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 used in the sense of
20 "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 (ALL) operation, block 2017, which can be considered one example of an execution resource) on this field e.g. increment or decrement
25 the condition match and the like. The CAM arrays are programmed 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
30 traversed that can then be used by other engines of the IP processor 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
35 signature of the classification path traversed or the like. The classification tag fields are chosen based on processor implementation and functionality. 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 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 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, 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 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 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 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 show 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 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 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 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 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 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 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-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 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 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 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 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 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 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 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, 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
5 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 involving
10 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
15 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
20 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
25 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 Fig.
30 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
35 receives the packets that need to be sent on to the network media independent interface 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
5 destination. The output queue controller may assign packets onto their respective quality of 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.

10 Fig. 26 illustrates the storage flow controller /RDMA controller block, shown generally at 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
15 command queues, new and active, blocks 2611 and 2610, and completion queue, block 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
20 include information like RDMA region and the access keys, and virtual address translation 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
25 storage flow controller schedules the commands for execution and also provides the 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
30 exist, to the scheduler for initiating a new connection. The scheduler 1702, uses the control 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
35 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 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 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, 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 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 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 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 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-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 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 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 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 encryption/decryption services. It is possible to use the message authentication service and the encryption/decryption services independent of each other.

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 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 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 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

5 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 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

10 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 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

15 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 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

20 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 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.

25 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 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

30 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 ALL) 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

35 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 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 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 the following major pipe stages:

1. Receive Pipe Stage of block 3201, with major steps illustrated in block 3207:

10 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 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.

15 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 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 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.

20 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 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 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 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 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 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 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 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. 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 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
5 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 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
10 purposes.

Fig. 33 illustrates write command operation between an initiator and a target. The Initiator 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
15 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 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
20 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 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 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,
30 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
35 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
5 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
10 policies that may be setup.

Fig. 35 illustrates the READ data transfer using RDMA mechanism between and 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,
15 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
20 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
25 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
30 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
35 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
5 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
10 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
15 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
20 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
25 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
30 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);

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 queue;
- 5 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
12. The packet is transferred to the output network media interface block which transmits the data packet to the destination.

10 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 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 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 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;
- 5 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;
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. It also requests the host to allocate appropriate buffer space for unsolicited data transfers
5 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 and queued to the input queue after layer 2 processing;

10 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;

12. The controller gets the ACK packet;

15 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.

The session establishment may also involve the login phase, which is not illustrated in the
20 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
25 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;
 - 5 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
10 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
15 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
20 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
25 the classification tag. The flow illustrated is for unencrypted packet and hence the security engine is not exercised;
 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);

15. The session entry is returned to the requesting packet processor;

5 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 descriptor or the memory descriptor list that may be used to perform the DMA of this packet into the host allocated destination buffer; and

10 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 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;

15 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

20 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;

25 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 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

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

- 5 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 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.
- 10 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 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
- 15 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 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
- 20 for applications. An efficient hardware implementation that can enable deployment of software driven security services is required to address the issues outlined above. The 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
- 25 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.

- Fig. 53 illustrates a traditional enterprise network with perimeter firewall. This figure
- 30 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,

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.

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

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 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 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. 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

5 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

10 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

15 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

20 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

25 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

30 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 reports

35 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 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 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 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 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 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 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

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 the trusted computing environment needs of the next generation computing systems.

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 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 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.

The IP processor 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
5 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
10 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.

I claim:

1. A security system comprising a network,
said network comprising one or more networked systems of one or more types,
5 a plurality of said one or more networked systems comprising a hardware processor providing transport layer protocol processing,
said hardware processor comprising
a protocol processing engine to do transport layer protocol processing; or
a programmable rule processing engine to analyze network traffic for rule
10 matching or taking actions on matched rules or a combination thereof; or
a security processing engine to do encryption, decryption, authorization or authentication or a combination thereof using standard or proprietary security protocols; or
a packet classification engine to classify the network traffic; or
15 a packet processing engine to perform packet processing tasks; or a combination of any of the foregoing,
said security system providing multiple protocol layer security in said network.
2. A security system for a storage area network,
said storage area network comprising one or more networked systems of one
20 or more types,
said security system comprising a set of systems from said one or more networked systems,
a plurality of said set of systems comprising a hardware processor providing transport layer protocol processing,
25 said hardware processor comprising
a storage protocol processing engine to do protocol processing; or

a protocol processing engine to do transport layer protocol processing; or

a programmable rule processing engine to analyze storage area network traffic for rule matching or taking actions on matched rules or a combination thereof; or

5 a security processing engine to do encryption, decryption, authorization or authentication or a combination thereof using standard or proprietary security protocols; or

a packet classification engine to classify the storage area network traffic; or

10 a packet processing engine to perform packet processing tasks like header processing or deep packet processing or a combination thereof; or

a combination of any or the foregoing,

said security system providing multiple protocol layer security in said storage area network.

3. The security system of claim 1 further comprising:

15 a. at least one central manager for compiling and distributing security rules; and

b. at least one security policy driver to communicate with the central manager and to set up rules in said hardware processor on at least one of said plurality of said one or more networked systems to analyze and enforce security based on said rules.

20 4. The security system of claim 3 wherein the central manager comprises at least one of:

a. A Security Policy Developer Interface for entering security policy;

b. A Security Rules Compiler for compiling security policies into rules;

25 c. A Rules Distribution Engine to distributed rules to said plurality of said one or more networked systems;

d. A Security Policy Manager Interface to manage said plurality of said one or more networked systems;

e. A Security Monitoring Engine to monitor said network;

f. An event collection/management engine to manage said network and collect events or reports from at least one of said plurality of said one or more networked systems; or

5 g. a combination of any of the foregoing.

5. The security system of claim 3 wherein at least one of said networked systems provides security based on rules for

a. OSI protocol layer two to provide layer two or MAC layer filtering; or

10 or b. OSI protocol layer three to provide layer three or network layer filtering;

c. OSI protocol layer four to provide layer four or transport layer filtering;
or

d. OSI protocol layers five through seven to provide upper layer or application layer filtering; or

15 e. a combination of any of the foregoing.

6. The security system of claim 1 including security protocols comprising at least one of IPSEC, OPSEC, SSL, TLS, AES, DES, 3DES, SHA1, MD4, MD5, RSA, CHAP, Kerberos, a proprietary protocol, or a combination of any of the foregoing.

20 7. The security system of claim 3 wherein at least one of the at least one policy drivers executes on a processor of said hardware processor or on a host processor of at least one of said networked systems.

25 8. The security system of claim 1 including multiple protocol layer security that includes security functions performed at one or more protocol layers of the OSI stack to provide packet filtering, intrusion detection, denial of service attack detection, port scanning detection, virus scan, spam filtering, unauthorized access, or a combination of any of the foregoing.

9. A security system comprising a network,

said network comprising one or more networked systems of one or more

types,

a plurality of said one or more networked systems comprising a hardware processor providing remote direct memory access capability,

said hardware processor comprising

5 an RDMA mechanism for performing RDMA data transfer or

a protocol processing engine to do transport layer protocol processing; or

a programmable rule processing engine to analyze network traffic for rule matching or taking actions on matched rules or a combination thereof; or

10 a security processing engine to do encryption, decryption, authorization or authentication or a combination thereof using standard or proprietary security protocols; or

a packet classification engine to classify the network traffic; or

a packet processing engine to perform packet processing tasks like header processing or deep packet processing or a combination thereof; or

15 a combination of any of the foregoing,

said security system providing multiple protocol layer security in said network.

10. The security system of claim 9 where said hardware processor provides a transport layer remote direct memory access capability.

11. The security system of claim 9 further comprising:

20 a. at least one central manager for compiling and distributing security rules; and

b. at least one security policy driver to communicate with the central manager to set up rules in said hardware processor on at least one of said plurality of said one or more networked systems to analyze and enforce security based on said rules.

25 12. The security system of claim 11 wherein the central manager comprises at least one of:

a. A Security Policy Developer Interface for entering security policy;
b. A Security Rules Compiler for compiling security policies into rules;
c. A Rules Distribution Engine to distribute rules to said plurality of said one or more networked systems

5 d. A Security Policy Manager Interface to manage said plurality of said one or more networked systems;

e. A Security Monitoring Engine to monitor said network;

f. An event collection/management engine to manage said network and collect events or reports from said plurality of said one or more networked systems; or

10 g. a combination of any of the foregoing.

13. The security system of claim 11 wherein at least one of said networked systems provides security based on rules for

a. OSI protocol layer two to provide layer two or MAC layer filtering; or

15 b. OSI protocol layer three to provide layer three or network layer filtering; or

c. OSI protocol layer four to provide layer four or transport layer filtering;

or

d. OSI protocol layers five through seven to provide upper layer or application layer filtering; or

20 e. a combination of any of the foregoing.

14. The security system of claim 9 including security protocols comprising at least one of IPSEC, OPSEC, SSL, TLS, AES, DES, 3DES, SHA1, MD4, MD5, RSA, CHAP, Kerberos, a proprietary protocol or a combination of any of the foregoing.

15. The security system of claim 11 wherein at least one of the at least one policy driver that executes on a processor of said hardware processor or on a host processor of at least one of said networked systems.

16. The security system of claim 9 including multiple protocol layer security that

includes security functions performed at one or more protocol layers of the OSI stack to provide packet filtering, intrusion detection, denial of service attack detection, port scanning detection, virus scan, spam filtering, unauthorized access, or a combination of any of the foregoing.

5 17. The combination of claim 1 wherein said one or more networked systems comprises a blade server, thin server, media server, streaming media server, appliance server, Unix server, Linux server, Windows or Windows derivative server, AIX server, clustered server, database server, grid computing server, VOIP server, wireless gateway server, security server, file server, network attached storage server, game server, router,
10 switch, wireless access point, workstation, desktop computer, notebook computer, laptop computer, utility computing system or gateway device or a combination of any of the foregoing.

 18. The combination of claim 9 wherein said one or more networked systems comprises a blade server, thin server, media server, streaming media server, appliance
15 server, Unix server, Linux server, Windows or Windows derivative server, AIX server, clustered server, database server, grid computing server, VOIP server, wireless gateway server, security server, file server, network attached storage server, game server, router, switch, wireless access point, workstation, desktop computer, notebook computer, laptop computer, utility computing system or gateway device or a combination of any of the
20 foregoing.

 19. The security system of claim 1 wherein said packet processing steps include header processing or deep packet processing or a combination thereof.

 20. The security system of claim 2 further comprising:
 a. at least one central manager for compiling and distributing storage
25 area network security rules; and
 b. at least one security policy driver to communicate with the central manager to set up rules in said hardware processor on at least one of said plurality of said one or more networked systems to analyze and enforce storage area network security based on said rules.

30 21. The security system of claim 20 wherein the central manager comprises at least one of:

 a. A Security Policy Developer Interface for entering security policy;

- b. A Security Rules Compiler for compiling security policies into rules;
- c. A Rules Distribution Engine to distribute rules to the said plurality of said one or more networked systems
- d. A Security Policy Manager Interface to manage said plurality of said one or more networked systems;
- e. A Security Monitoring Engine to monitor said network;
- f. An event collection/management engine to manage said network and collect events or reports from said plurality of said one or more networked systems; or
- g. a combination of any of the foregoing.

10 22. The security system of claim 20 wherein at least one of said networked systems provides security based on rules for

- a. OSI protocol layer two to provide layer two or MAC layer filtering; or
- b. OSI protocol layer three to provide layer three or network layer filtering; or
- 15 c. OSI protocol layer four to provide layer four or transport layer filtering; or
- d. OSI protocol layers five through seven to provide upper layer or application layer filtering; or
- e. Storage protocol layer to provide storage protocol layer filtering; or
- 20 f. a combination of any of the foregoing.

23. The security system of claim 2 including security protocols comprising at least one of IPSEC, OPSEC, SSL, TLS, AES, DES, 3DES, SHA1, MD4, MD5, RSA, CHAP, Kerberos, a proprietary protocol or a combination of any of the foregoing.

24. The security system of claim 20 including a policy driver that executes on a processor of said hardware processor or on a host processor of at least one of said networked systems.

25. The security system of claim 2 including multiple protocol layer security that

includes security functions performed at one or more protocol layers of the OSI stack to provide packet filtering, intrusion detection, denial of service attack detection, port scanning detection, virus scan, spam filtering, unauthorized access, or a combination of any of the foregoing.

5 26. A security system for a network,

 said network comprising one or more networked systems of one or more types,

 said security system comprising a set of systems from said one or more networked systems,

10 a plurality of said set of systems comprising a hardware processor providing transport layer protocol processing,

 said hardware processor comprising a protocol processing engine to do transport layer protocol processing; or

15 a programmable rule processing engine to analyze network traffic for rule matching or taking actions on matched rules or a combination thereof; or

 a security processing engine to do encryption, decryption, authorization or authentication or a combination thereof using standard or proprietary security protocols; or

 a packet classification engine to classify the network traffic; or

20 a packet processing engine to perform packet processing tasks like header processing or deep packet processing or a combination thereof; or

 a combination of the foregoing,

 said security system providing multiple protocol layer security in said network.

25 27. A security system for a network comprising one or more networked systems, at least one of said networked systems having a hardware processor providing a protocol processing stack, said security system providing a secure operating environment for said protocol processing stack for trusted computing needs of one or more of said networked systems by providing a policy driver for setting up the hardware processor for security policy rules to be enforced by said hardware processor, and a central manager for compiling and

distributing said rules and monitoring the enforcement of said rules by said hardware processor.

Fig. 1 SCSI Architecture Layers

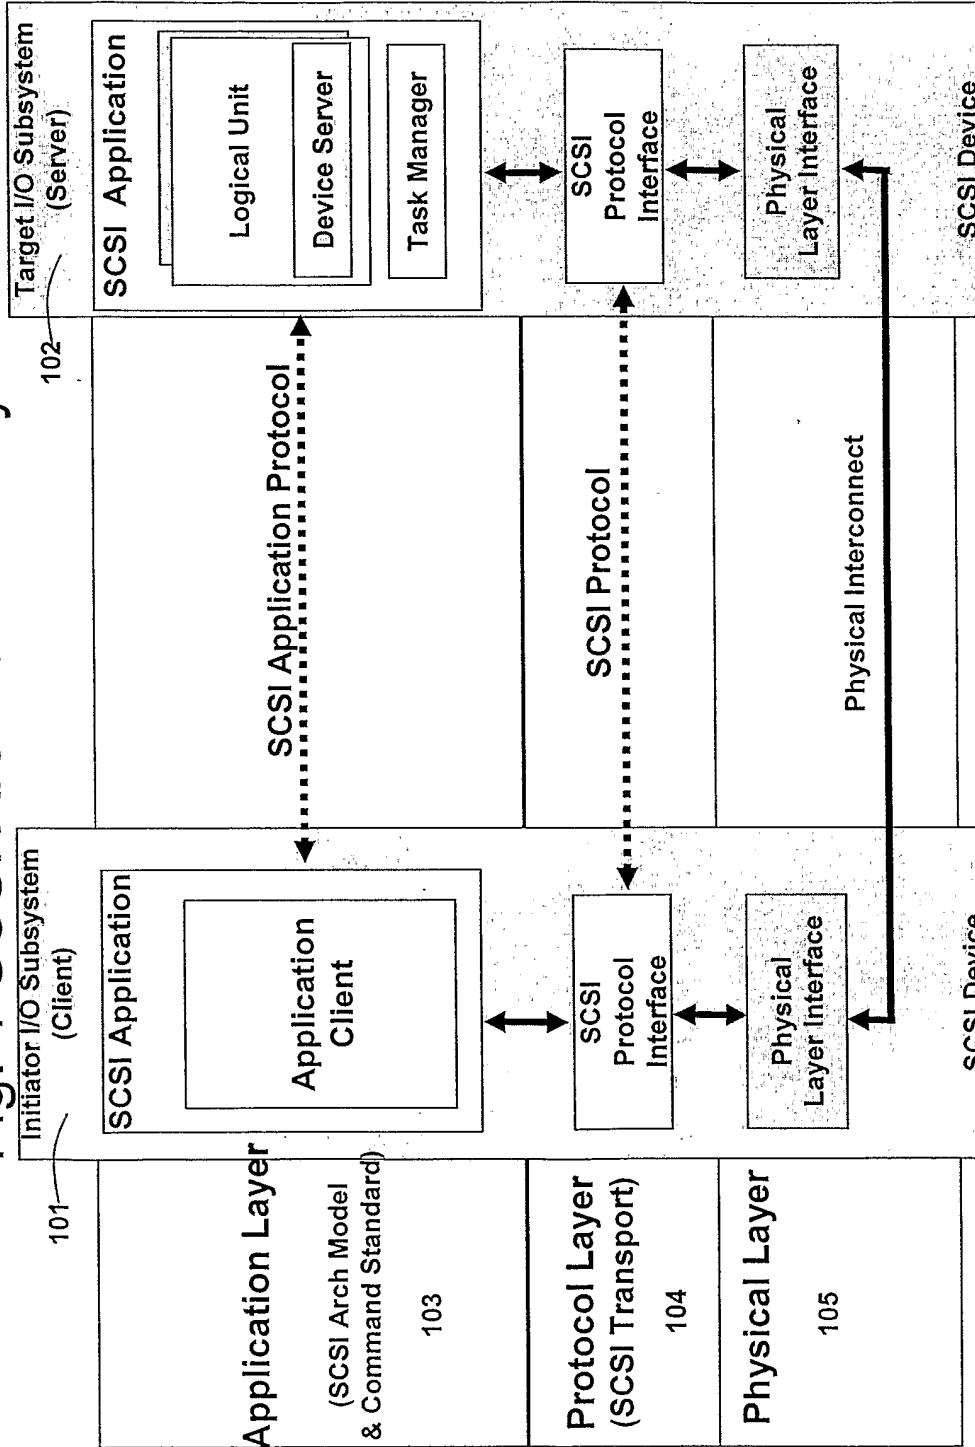


Fig. 2 SCSI Architecture Layers on iSCSI

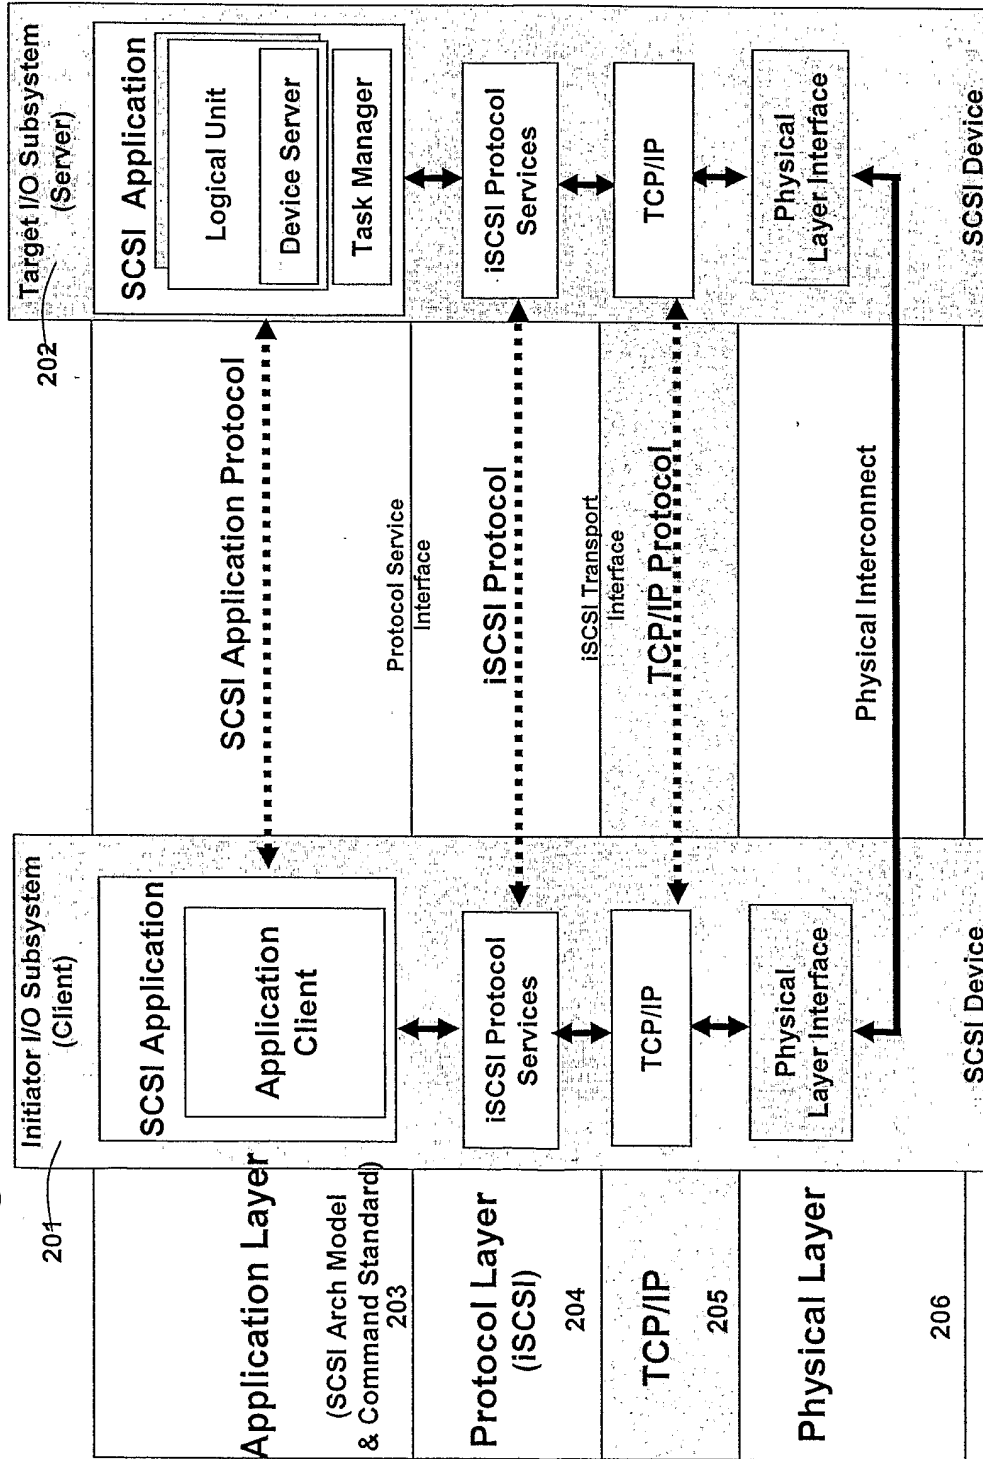
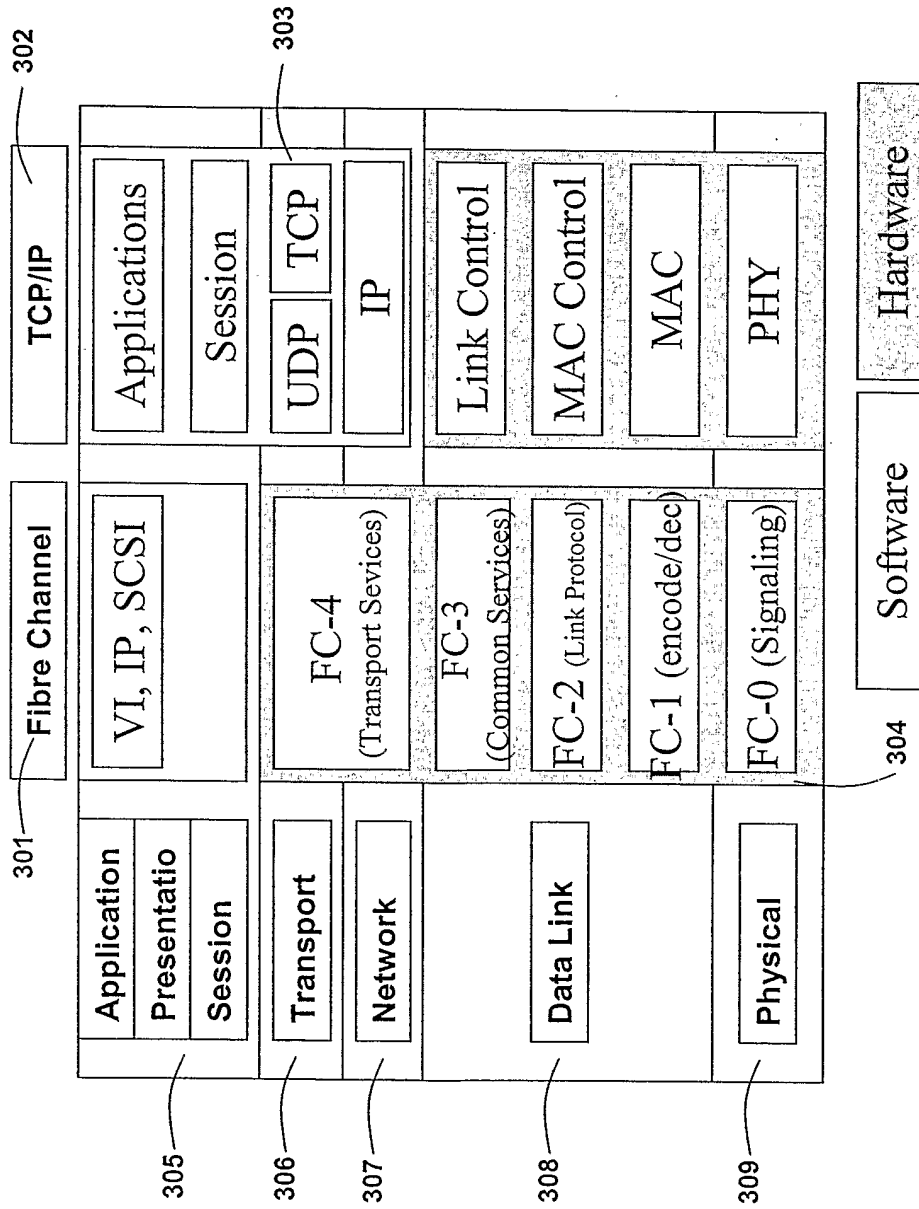


Fig. 3 OSI Stack



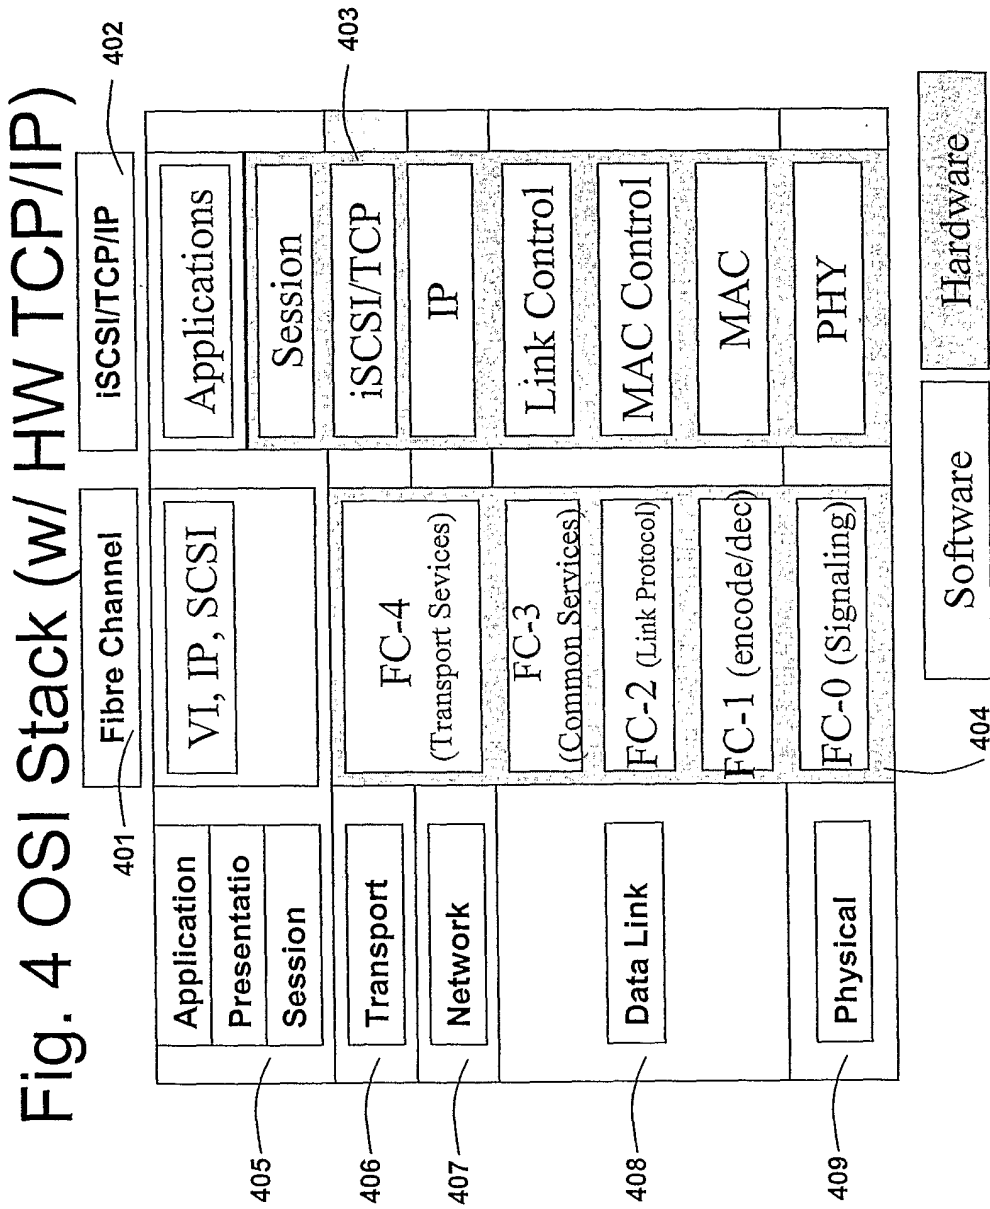


Fig. 5 Host Software Stack (iSCSI)




Fig. 6 S/W TCP Stack Data Transfer




Fig. 7 Remote Direct Memory Access

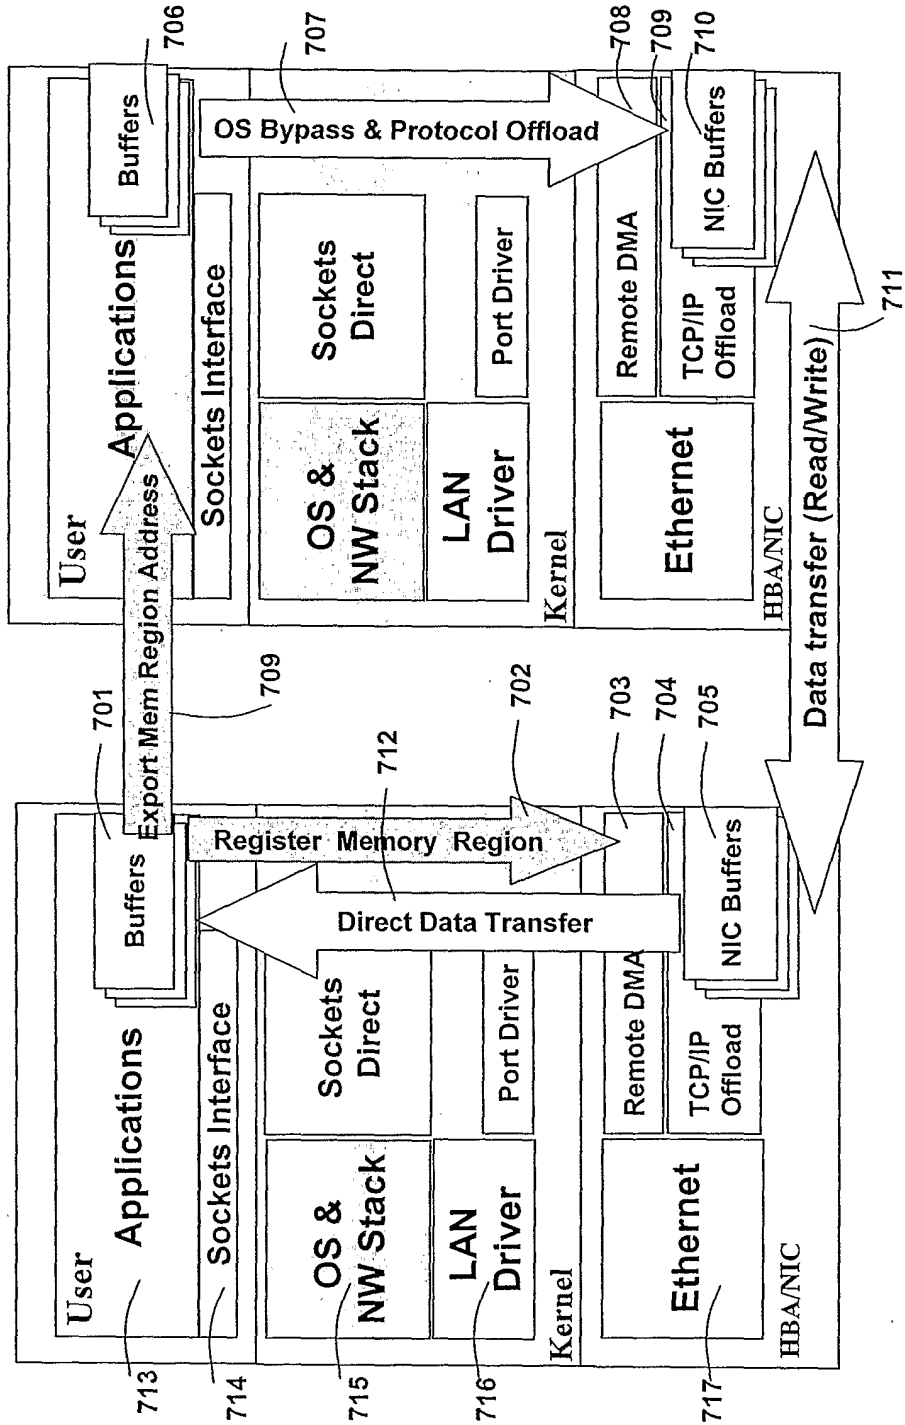


Fig. 8 Host Software Stack (SCSI)




Fig. 9 iSCSI Stack

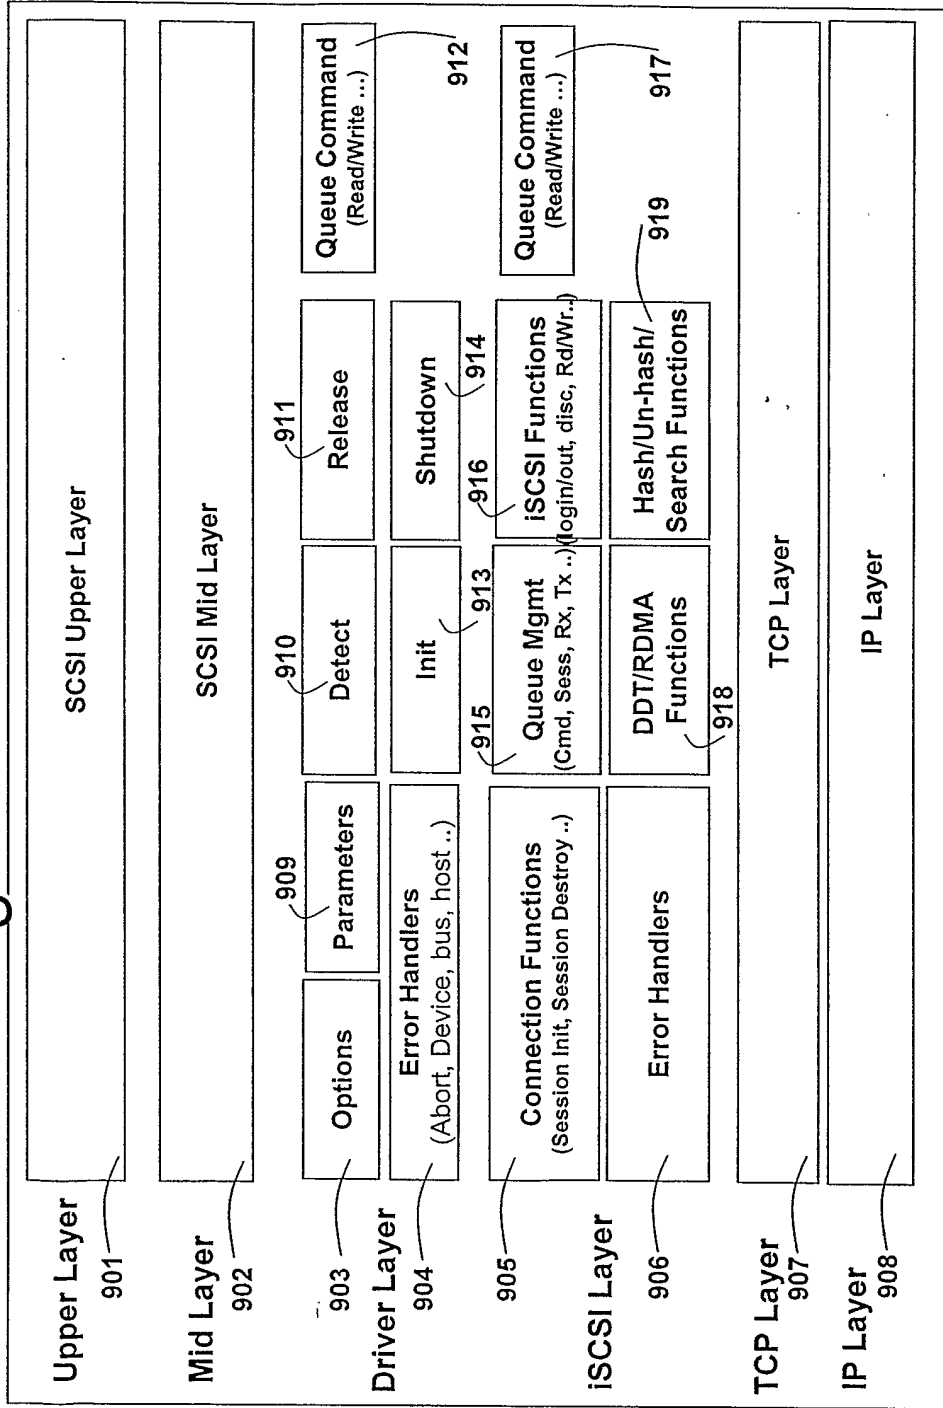


Fig. 10 TCP/IP Stack

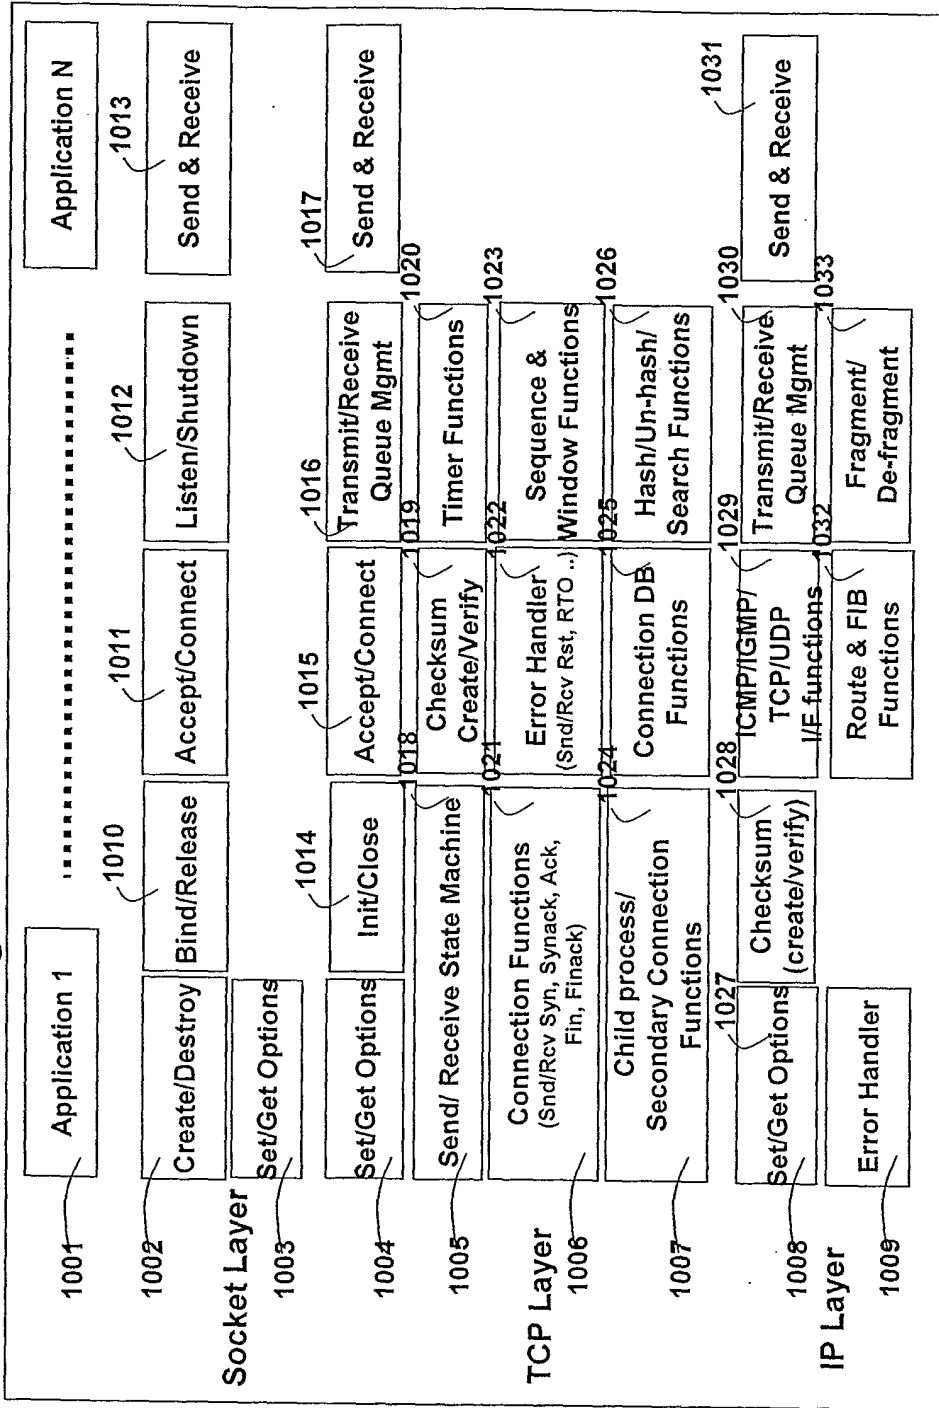


Fig. 11 iSCSI Data Flow




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_SQN |
| TCP_State | TCP_SND_SEQ# | TCP_SEQ# | TCP_ACK# | TCP_RCV_SEQ# |
| Fragment_PTR | Packet_send_PTR <small>(byte granular)</small> | Packet_Start_Time | Last_PKT_Time | TCP_Parameters <small>(slow start, retransmit ..)</small> |
| 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 rcv'ed) | 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

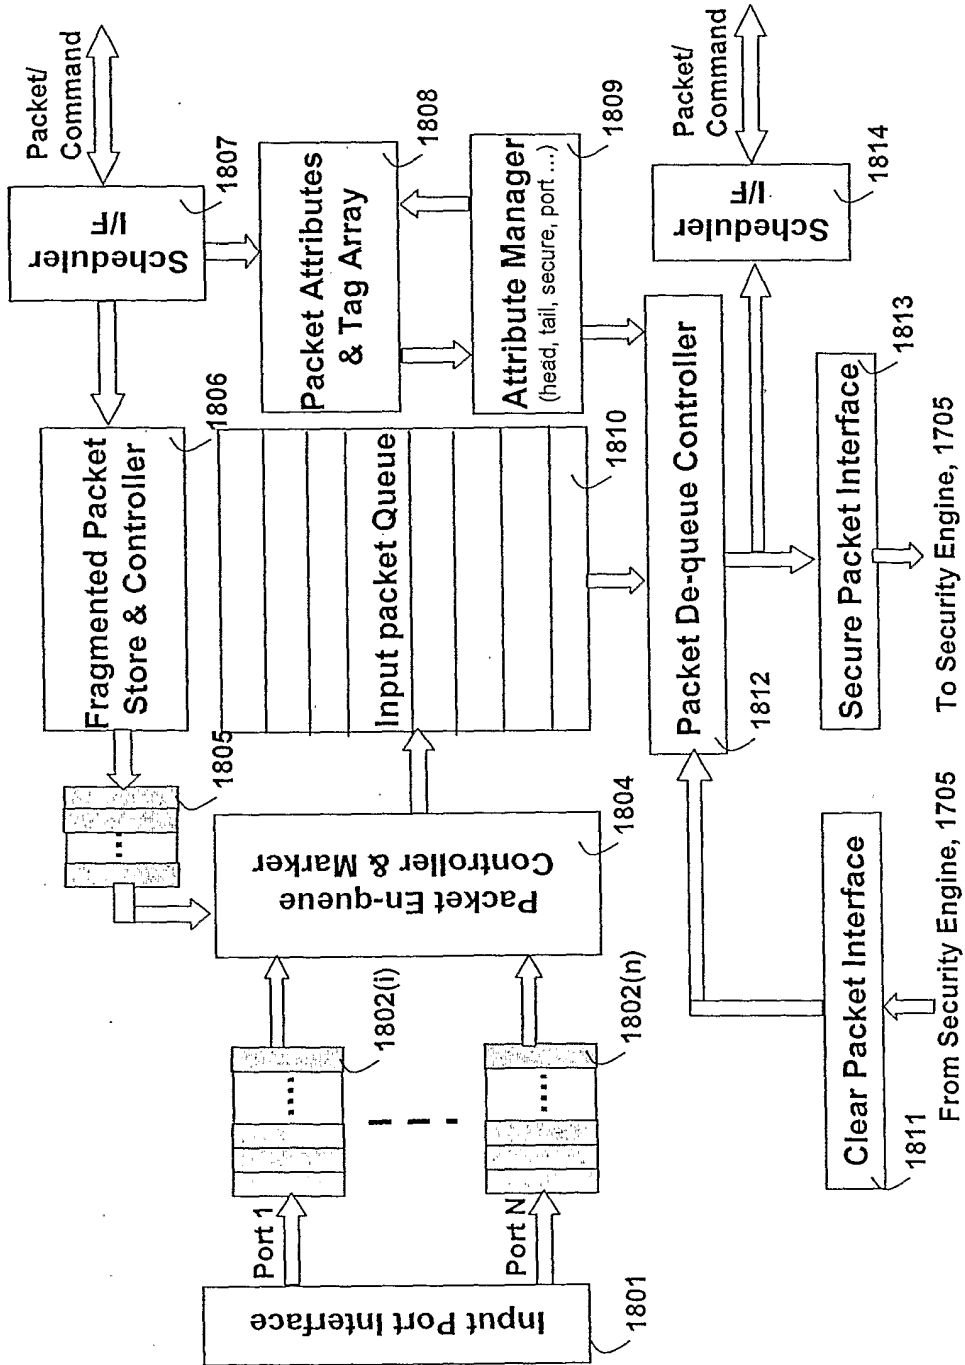


Fig. 20 Packet Classification/Policy Engine

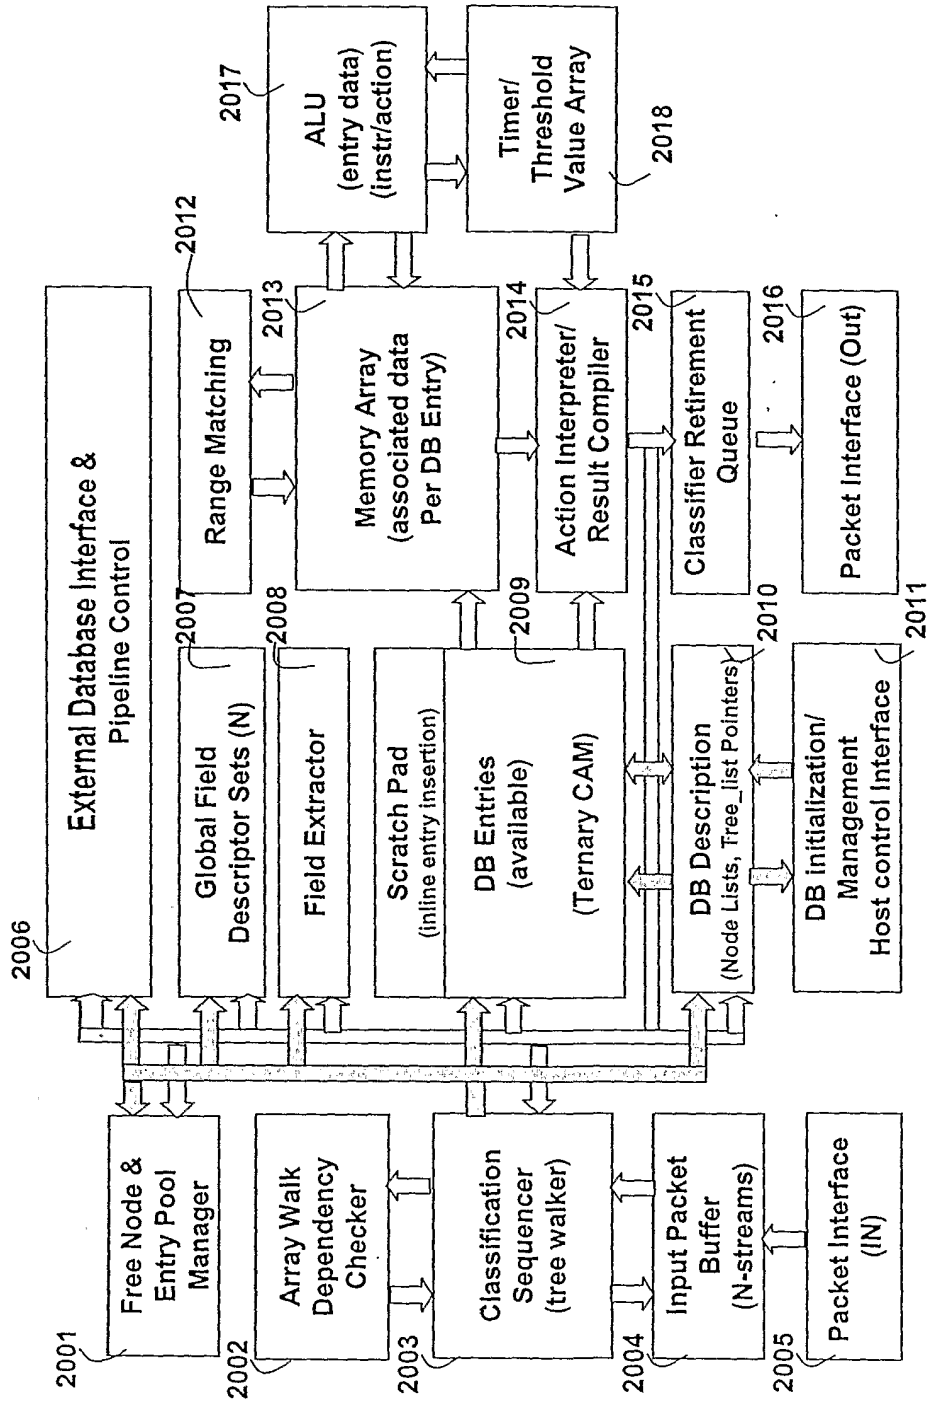


Fig. 21 SAN Packet Processor




Fig. 22 SAN Packet Processor




Fig. 23 TCP/IP Processor




Fig. 24 IP Storage Engine (iSCSI Processor)

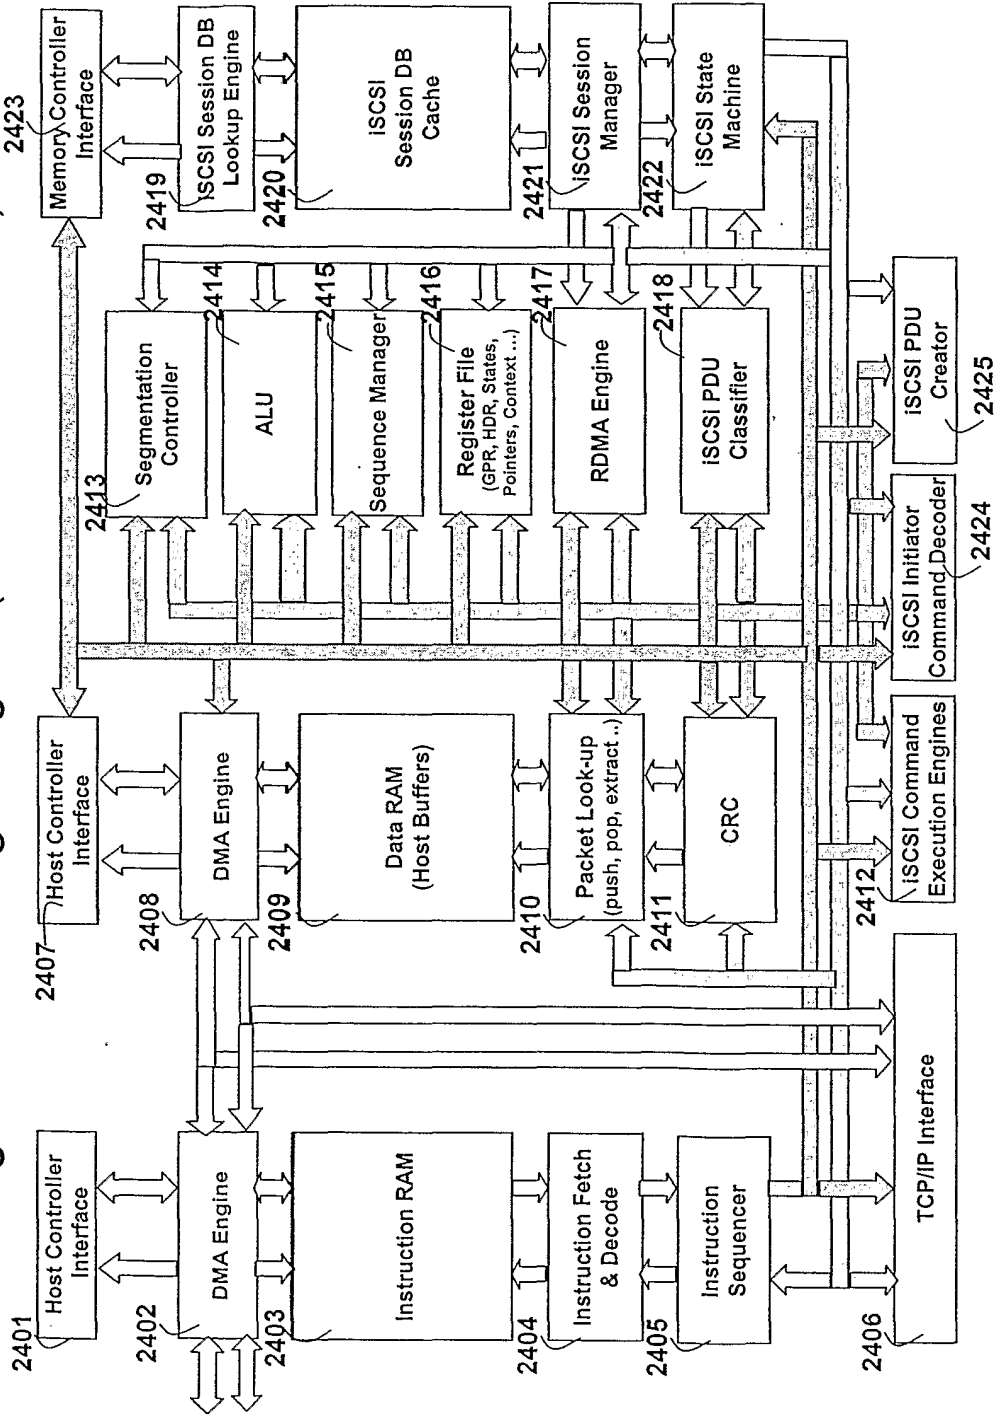


Fig. 25 Output Queue

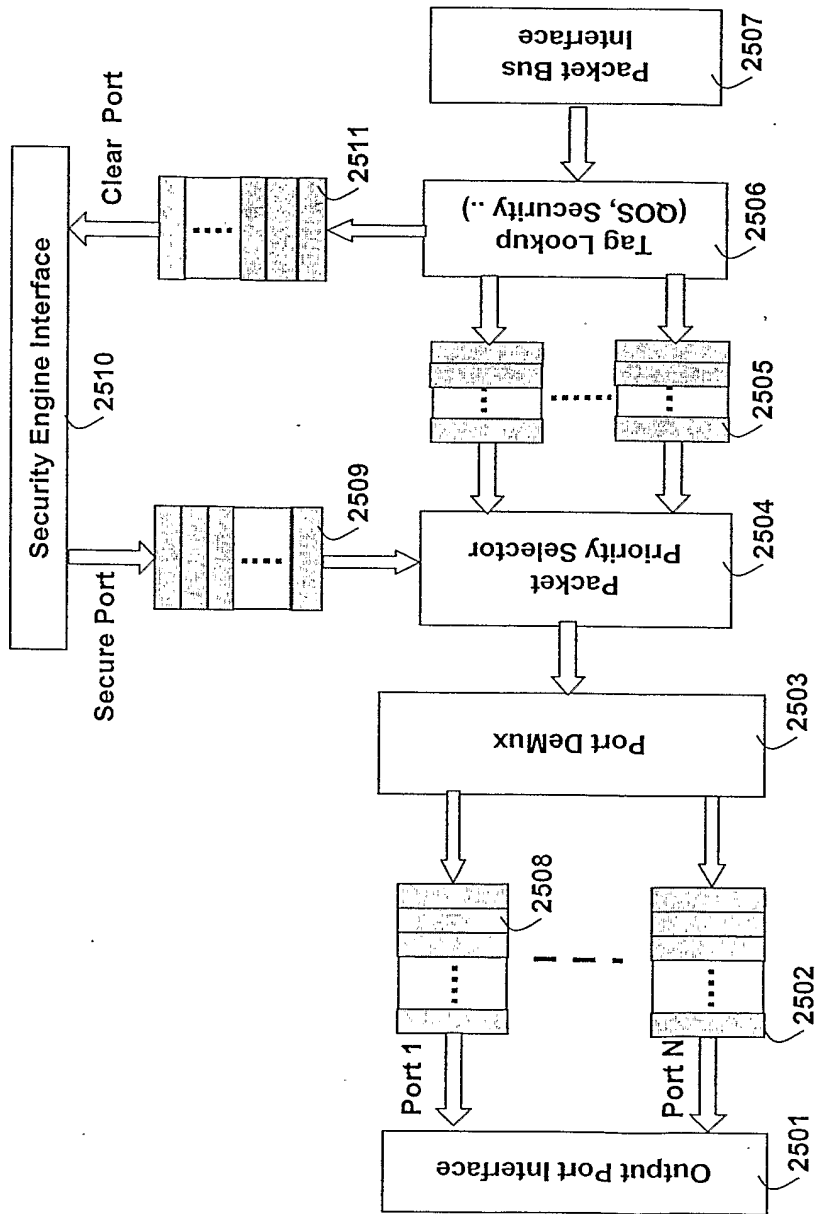
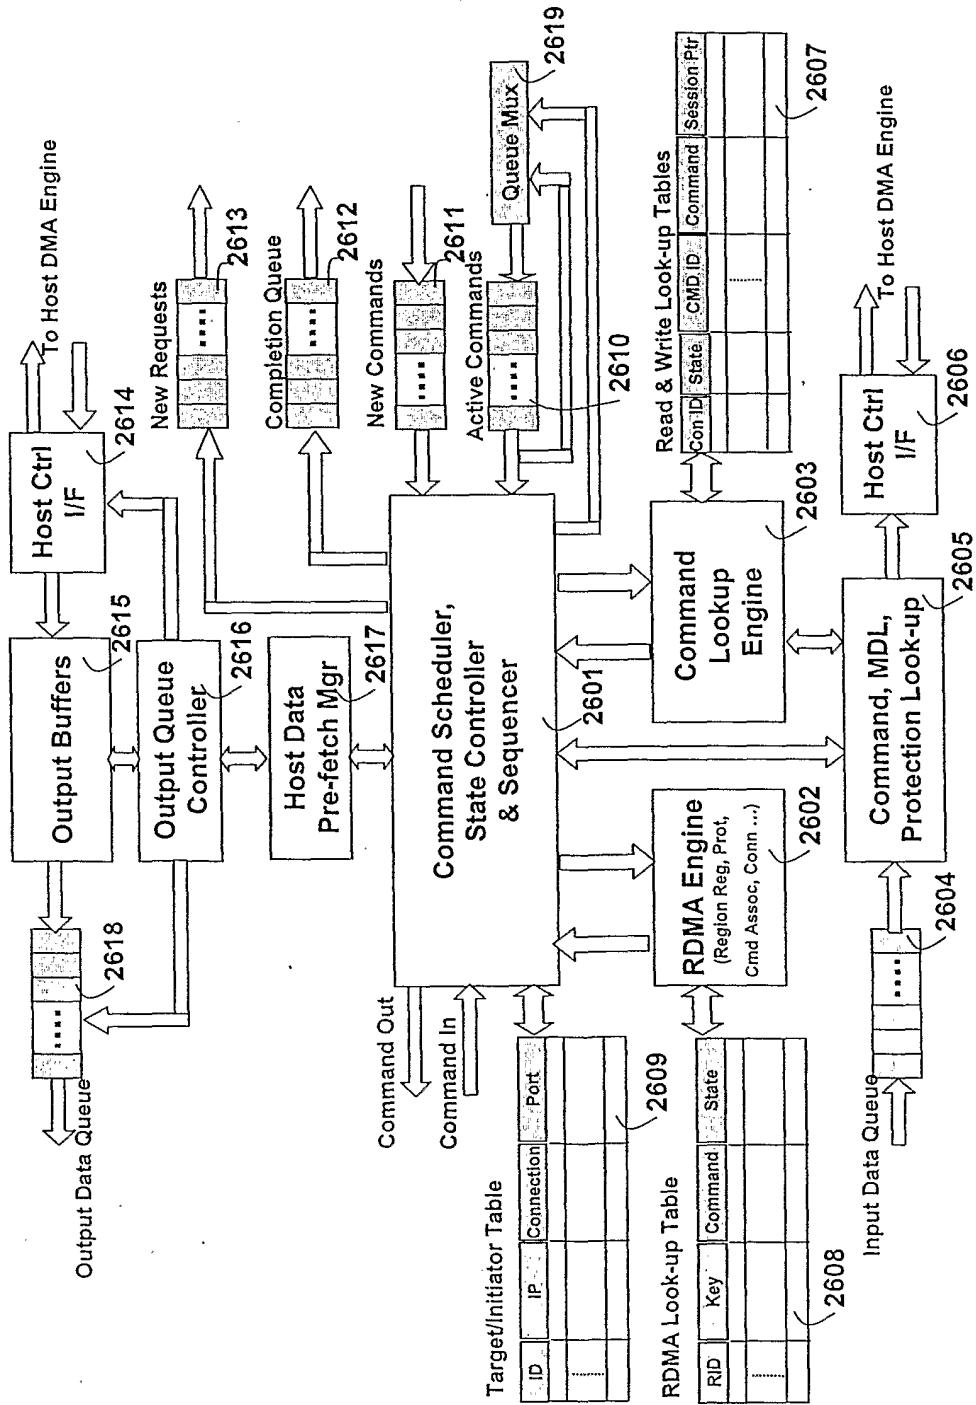


Fig. 26 Storage Flow & RDMA Controller



Target/Initiator Table

| ID | IP | Connection | Port |
|----|----|------------|------|
| | | | |
| | | | |
| | | | |

RDMA Look-up Table

| RID | Key | Command | State |
|-----|-----|---------|-------|
| | | | |
| | | | |
| | | | |

Read & Write Look-up Tables

| Con-ID | State | CMD-ID | Command | Session Ptr |
|--------|-------|--------|---------|-------------|
| | | | | |
| | | | | |
| | | | | |

Fig. 27 Host Interface Controller

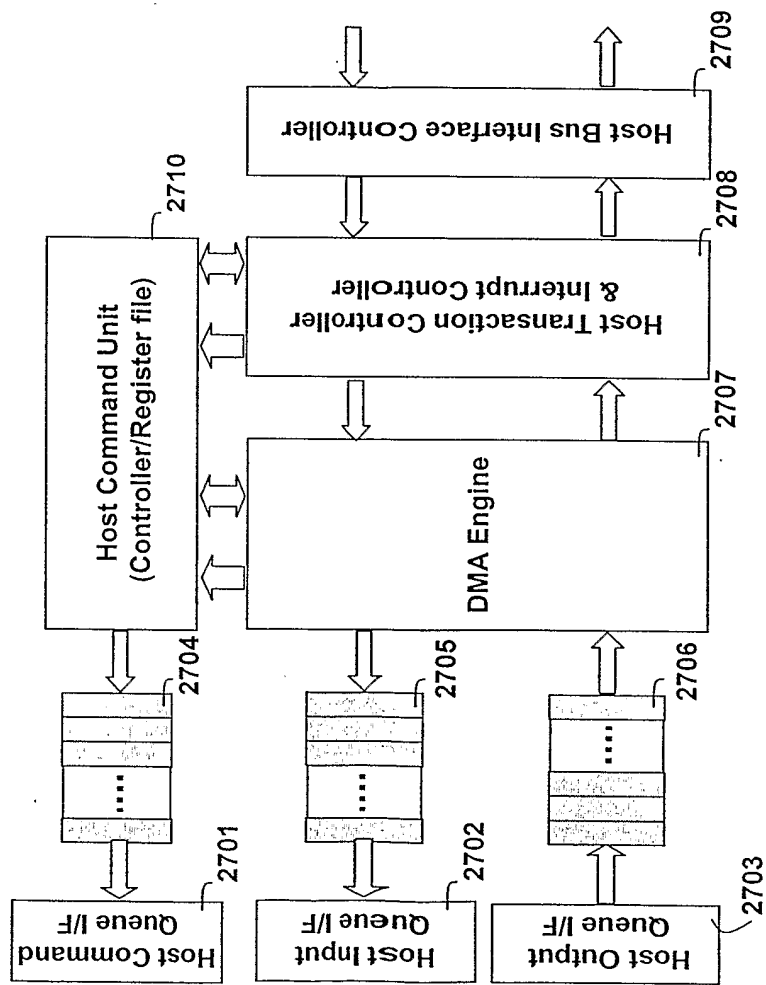


Fig. 28 Security Engine

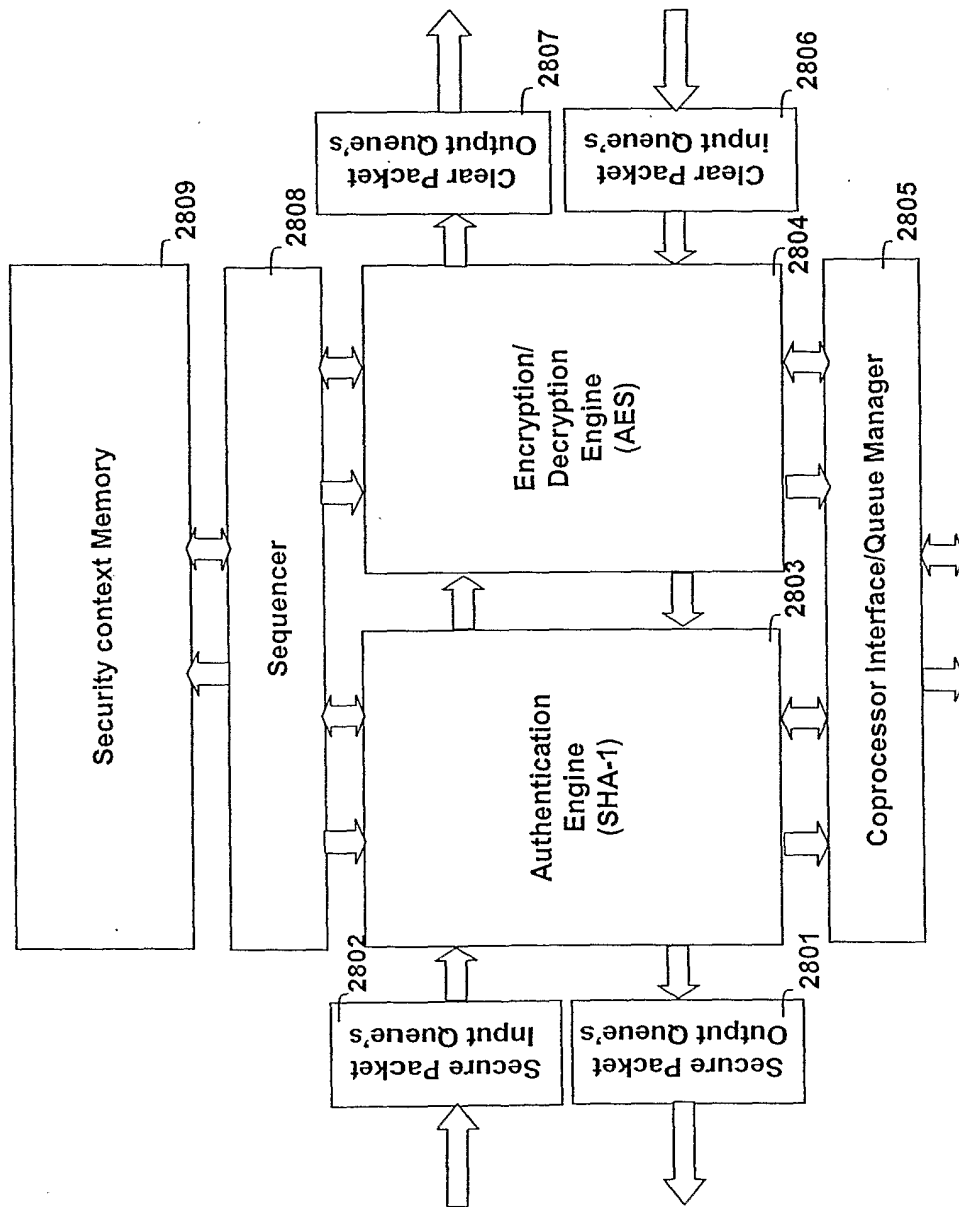


Fig. 29 IP Storage Session Cache & Memory Controller Complex

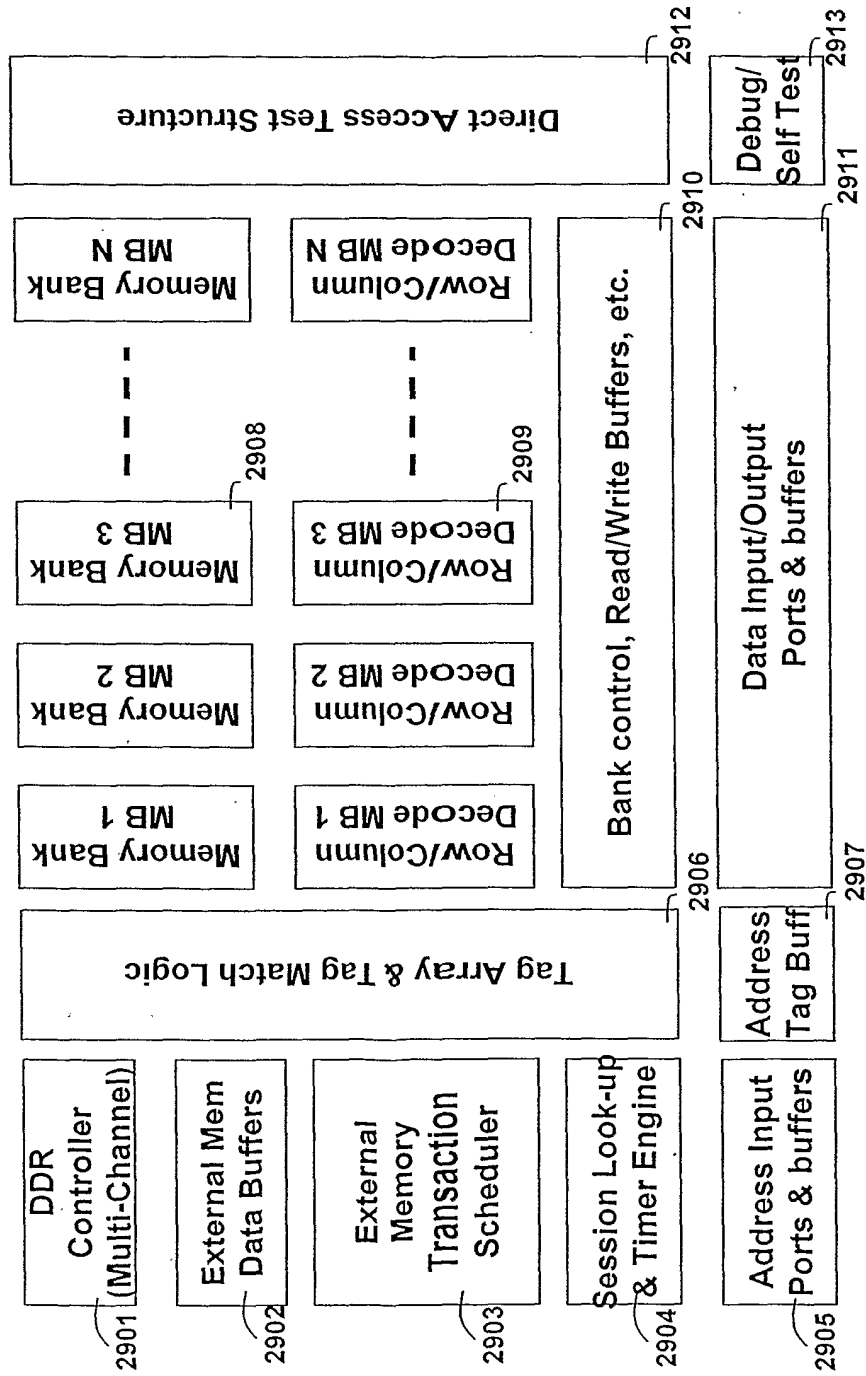


Fig. 30 Classification Engine Data Structure

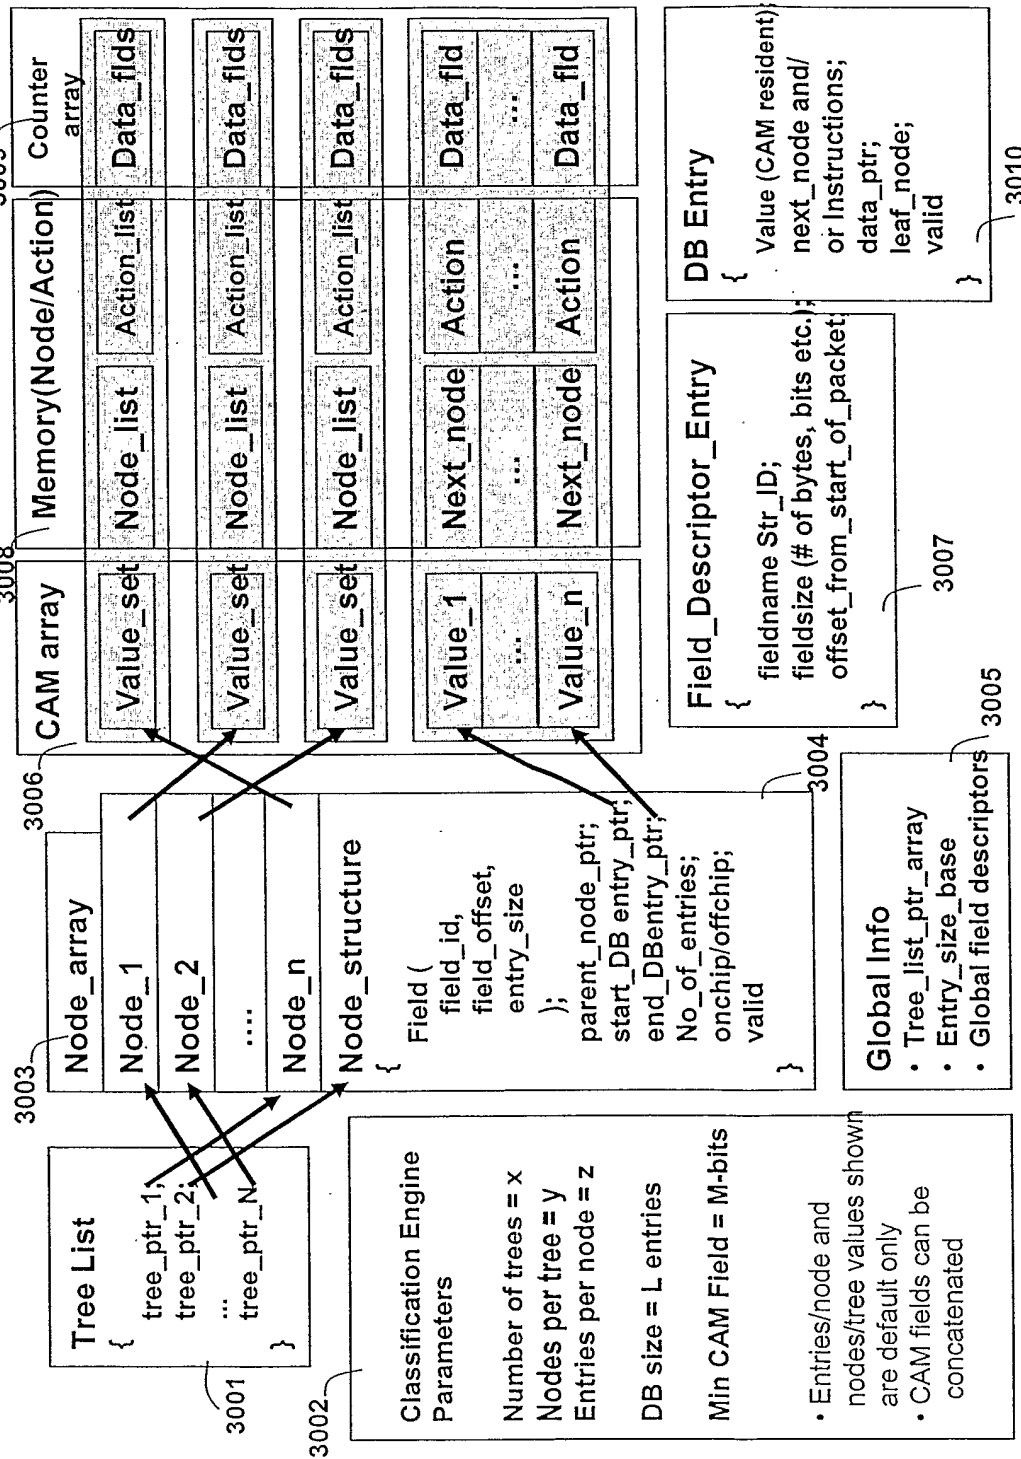


Fig. 31 Read Operation




Fig. 32 Read Data Packet Flow

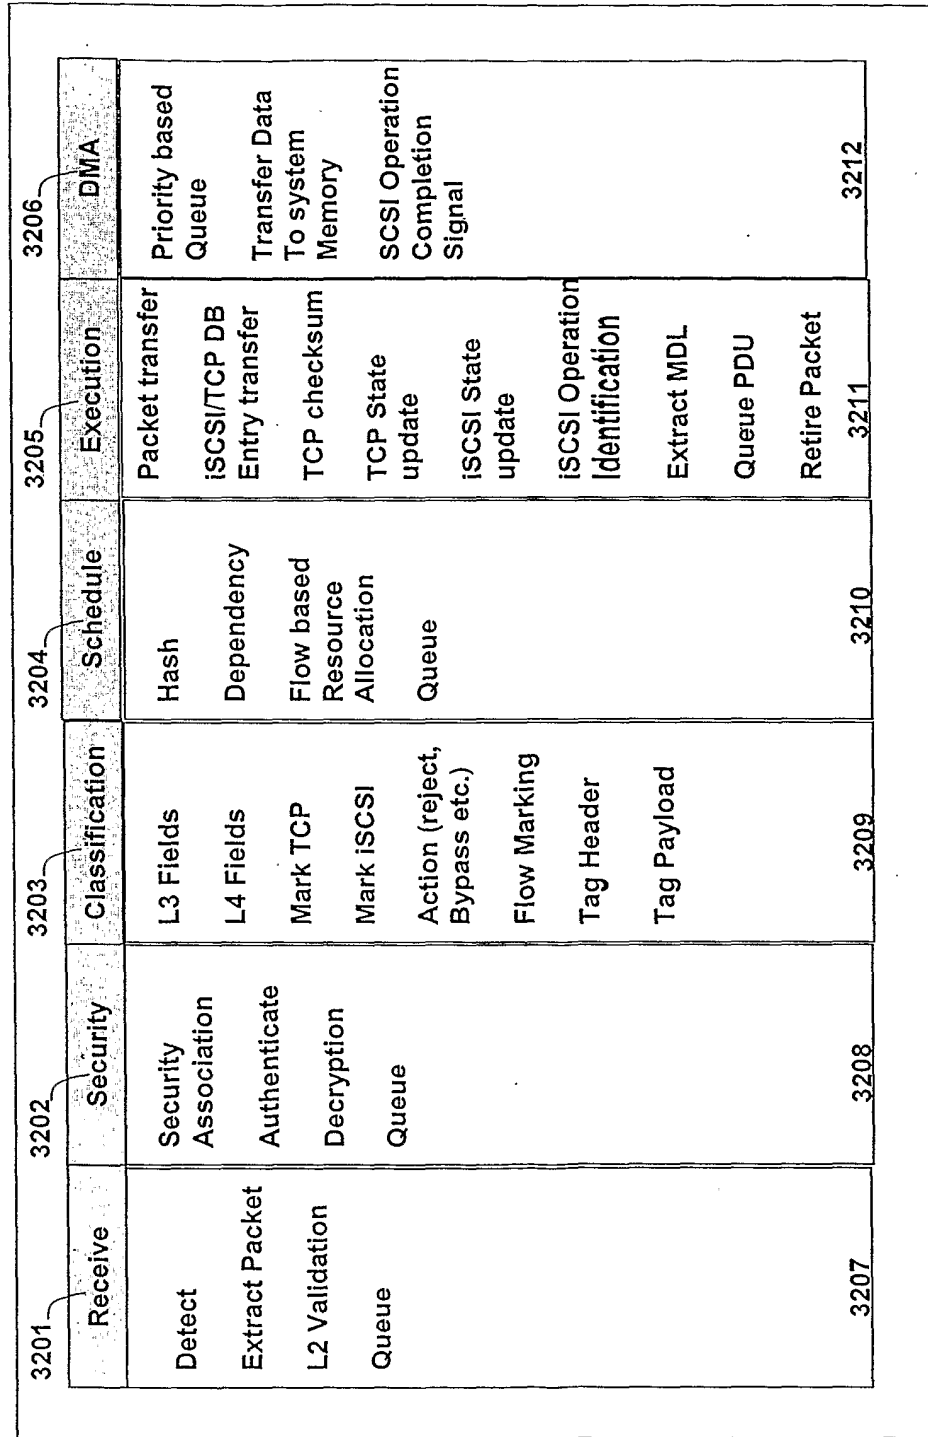


Fig. 33 Write Operation

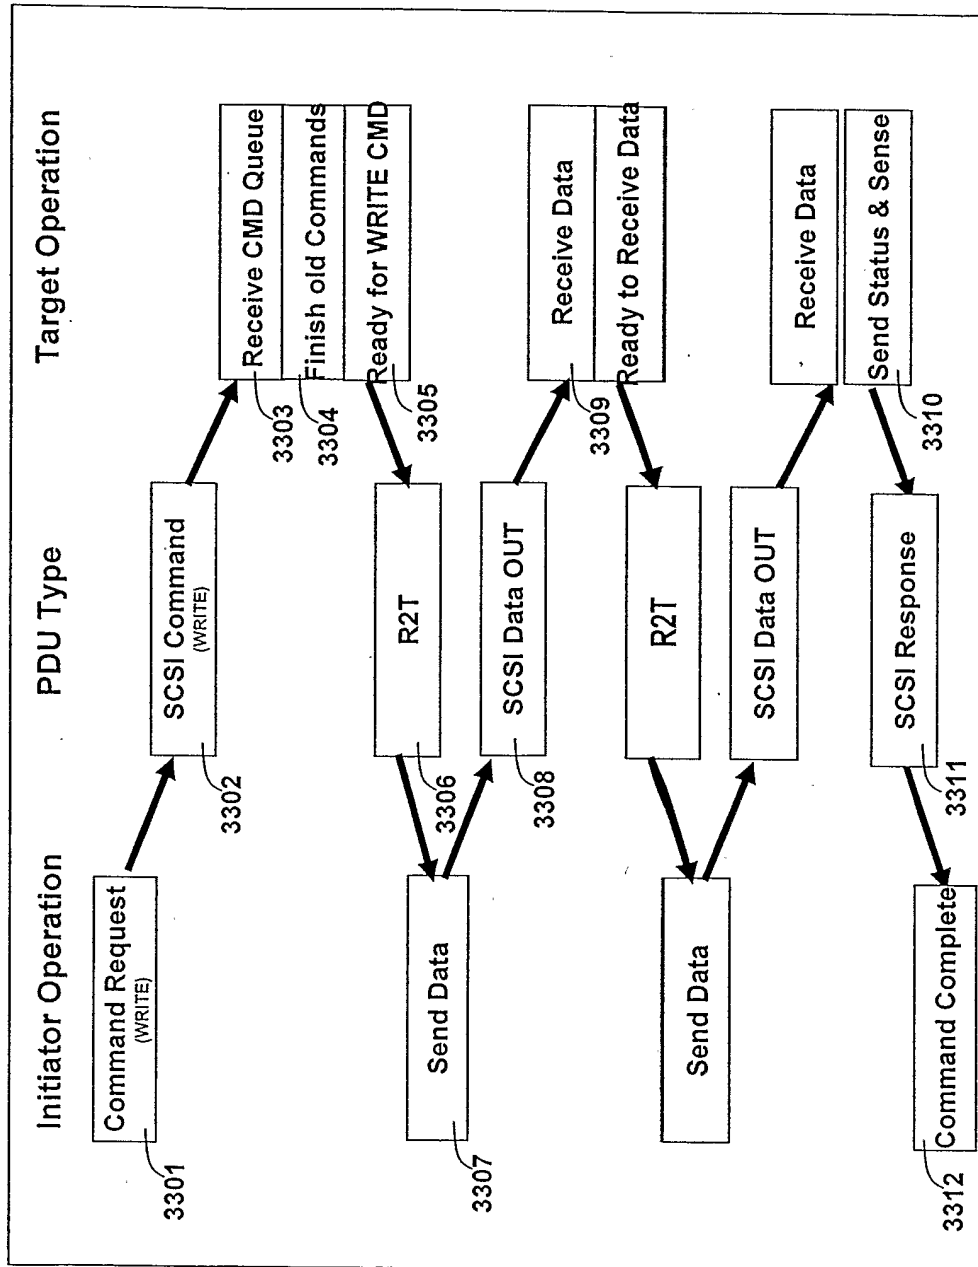


Fig. 34 Write Data Packet Flow

| 3401 | | 3402 | | 3403 | | 3404 | | 3405 | | 3406 | | 3407 | | 3408 | |
|-------------------|------------------------------|--------------------------------|--------------------------------|----------------|--------------------------------|----------------------|---------------------|----------------------------------|-----------------------------|-----------------------------|------------------|-----------------|----------|----------|--|
| Receive | Classification | Schedule | Execution | DMA | Execution | Security | Transmit | | | | | | | | |
| Detect | L3 Fields | Hash Key Lookup | Packet transfer | Priority Queue | Packet transfer | Security Association | Assemble L2 Headers | Transfer Data from system Memory | ISCSI/TCP DB Entry transfer | Encryption | L2 Headers Queue | Prepare Headers | Transmit | Checksum | |
| Extract Packet | L4 Fields | Dependency | TCP checksum | | TCP State update | Message Digest Queue | | | ISCSI/TCP DB Entry transfer | Prepare Packet | | | | | |
| L2 Validation | Mark TCP | Flow based Resource Allocation | TCP State update | | ISCSI State update | | | | TCP State update | ISCSI/TCP DB Entry transfer | | | | | |
| Queue | Mark ISCSI | Queue | ISCSI State update | | ISCSI Operation Identification | | | | ISCSI State update | | | | | | |
| | Action (reject, Bypass etc.) | | ISCSI Operation Identification | | Extract MDL | | | | Packet transfer | | | | | | |
| | Flow Marking | | Extract MDL | | Queue PDU | | | | Queue Packet | | | | | | |
| | Tag Header | | Start Write DMA | | Retire Packet | | | | | | | | | | |
| | Tag Payload | | | | | | | | | | | | | | |
| 3415 | 3416 | 3409 | 3410 | 3411 | 3410 | 3413 | 3414 | | | | | | | | |
| R2T Read Response | | | | | | | | Write Data Transfer | | | | | | | |

Fig. 35 iSCSI Read using RDMA Write

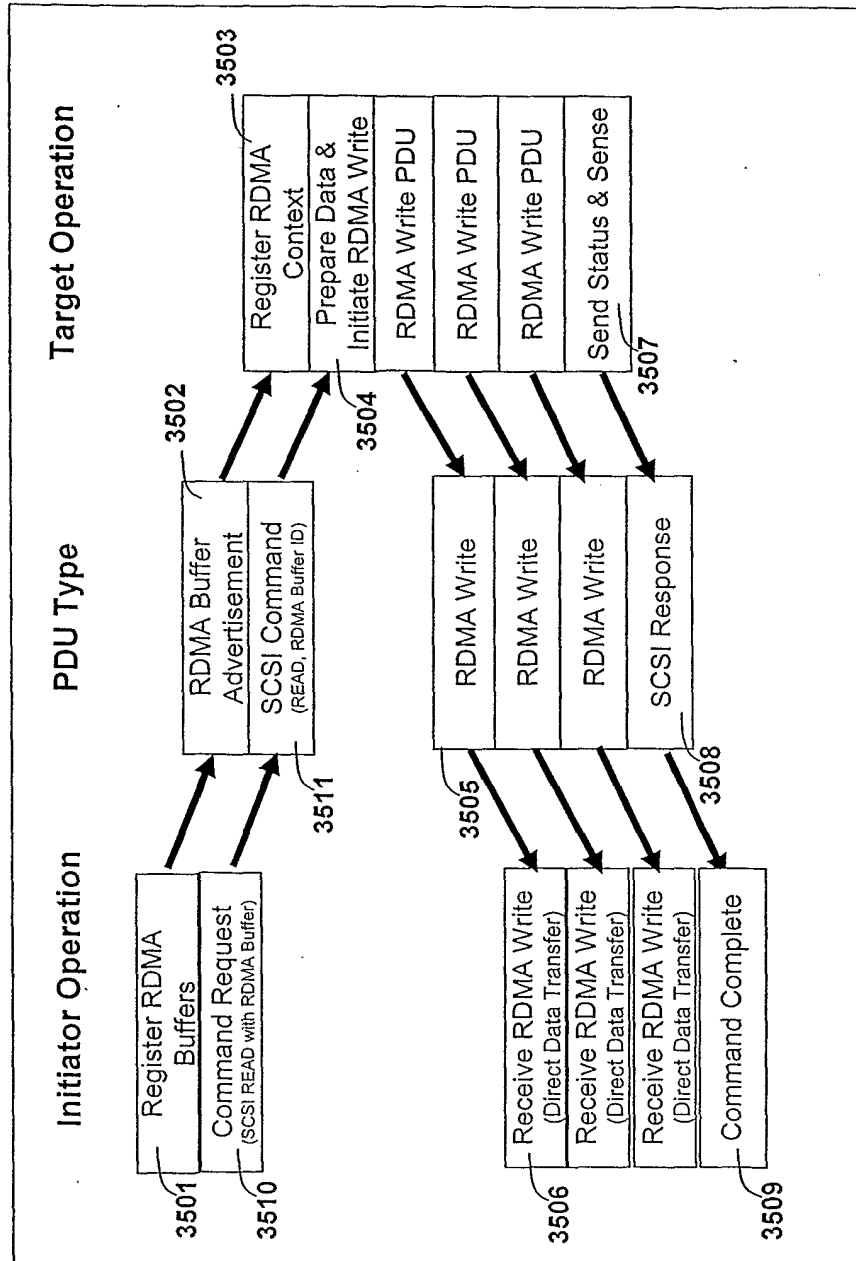


Fig. 36 Read Data Packet Flow (RDMA)

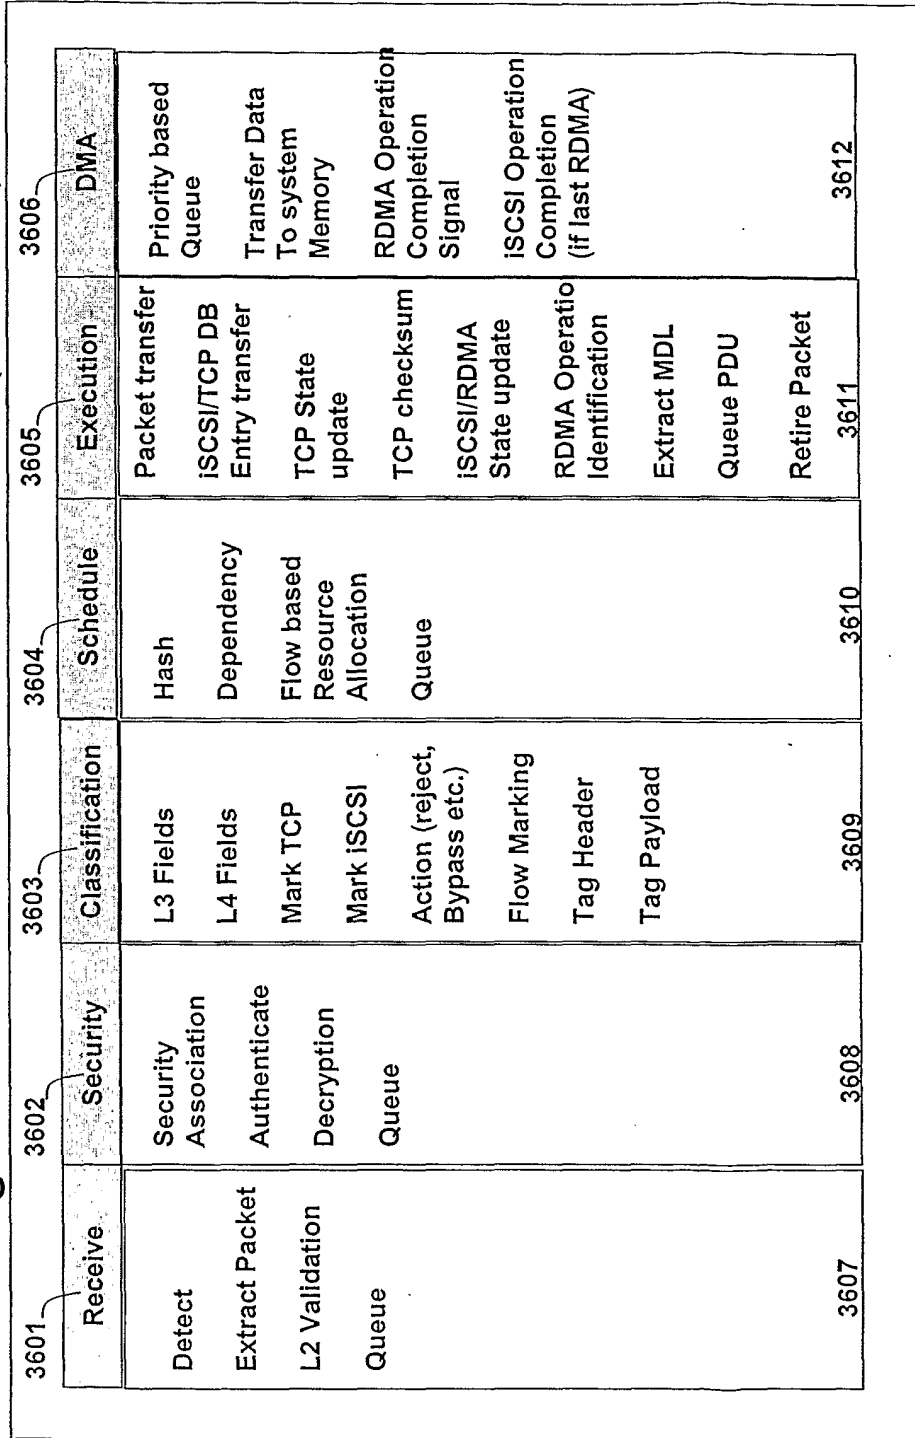


Fig. 37 iSCSI Write using RDMA Read

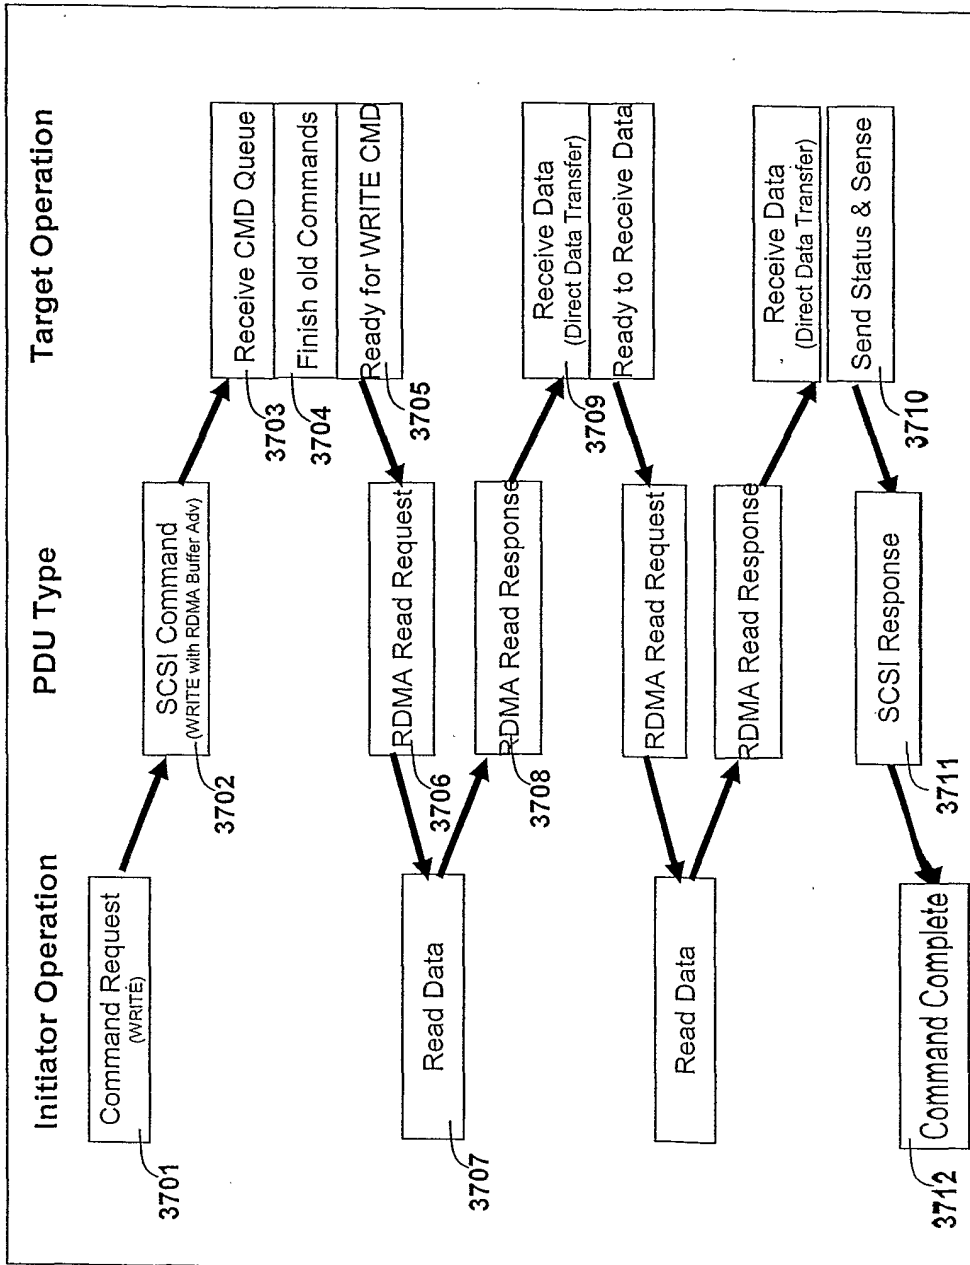


Fig. 38 Write Data Packet Flow (RDMA)

| 3801 | | 3802 | | 3803 | | 3804 | | 3805 | | 3806 | | 3807 | | 3808 | |
|-------------------|------------------------------|--------------------------------|--------------------------|----------------|-------------------------|----------------------|----------|-----------------------------------|----------------|------------|-----------------------------|----------------------------------|-----------|------------|------------|
| Receive | Classification | Schedule | Execution | DMA | Execution | Security | Transmit | Receive | Classification | Schedule | Execution | DMA | Execution | Security | Transmit |
| Detect | L3 Fields | Hash Key Lookup | Packet transfer | Priority Queue | Prepare Headers | Security Association | Assemble | Extract Packet | L4 Fields | Dependency | iSCSI/TCP DB Entry transfer | Transfer Data from system Memory | Checksum | Encryption | L2 Headers |
| L2 Validation | Mark TCP | Flow based Resource Allocation | TCP checksum | | Prepare Packet | Message Digest Queue | Queue | Queue | Mark iSCSI | Queue | iSCSI/TCP DB Entry transfer | | | | |
| Queue | Action (reject, Bypass etc.) | Queue | TCP State update | | iSCSI/RDMA State update | Queue | Transmit | | Flow Marking | | TCP State update | | | | |
| | Tag Header | | iSCSI/RDMA State update | | RDMA Op Identification | | | | Tag Payload | | iSCSI/RDMA State update | | | | |
| | | | Extract RID Access Check | | Packet transfer | | | | | | Queue Packet | | | | |
| | | | Start Write RDMA | | Queue Packet | | | | | | | | | | |
| | | | Retire Packet | | | | | | | | | | | | |
| 3815 | 3816 | 3809 | 3810 | 3811 | 3812 | 3813 | 3814 | | | | | | | | |
| RDMA Read Request | | | | | | | | RDMA Read Response Packet (Write) | | | | | | | |

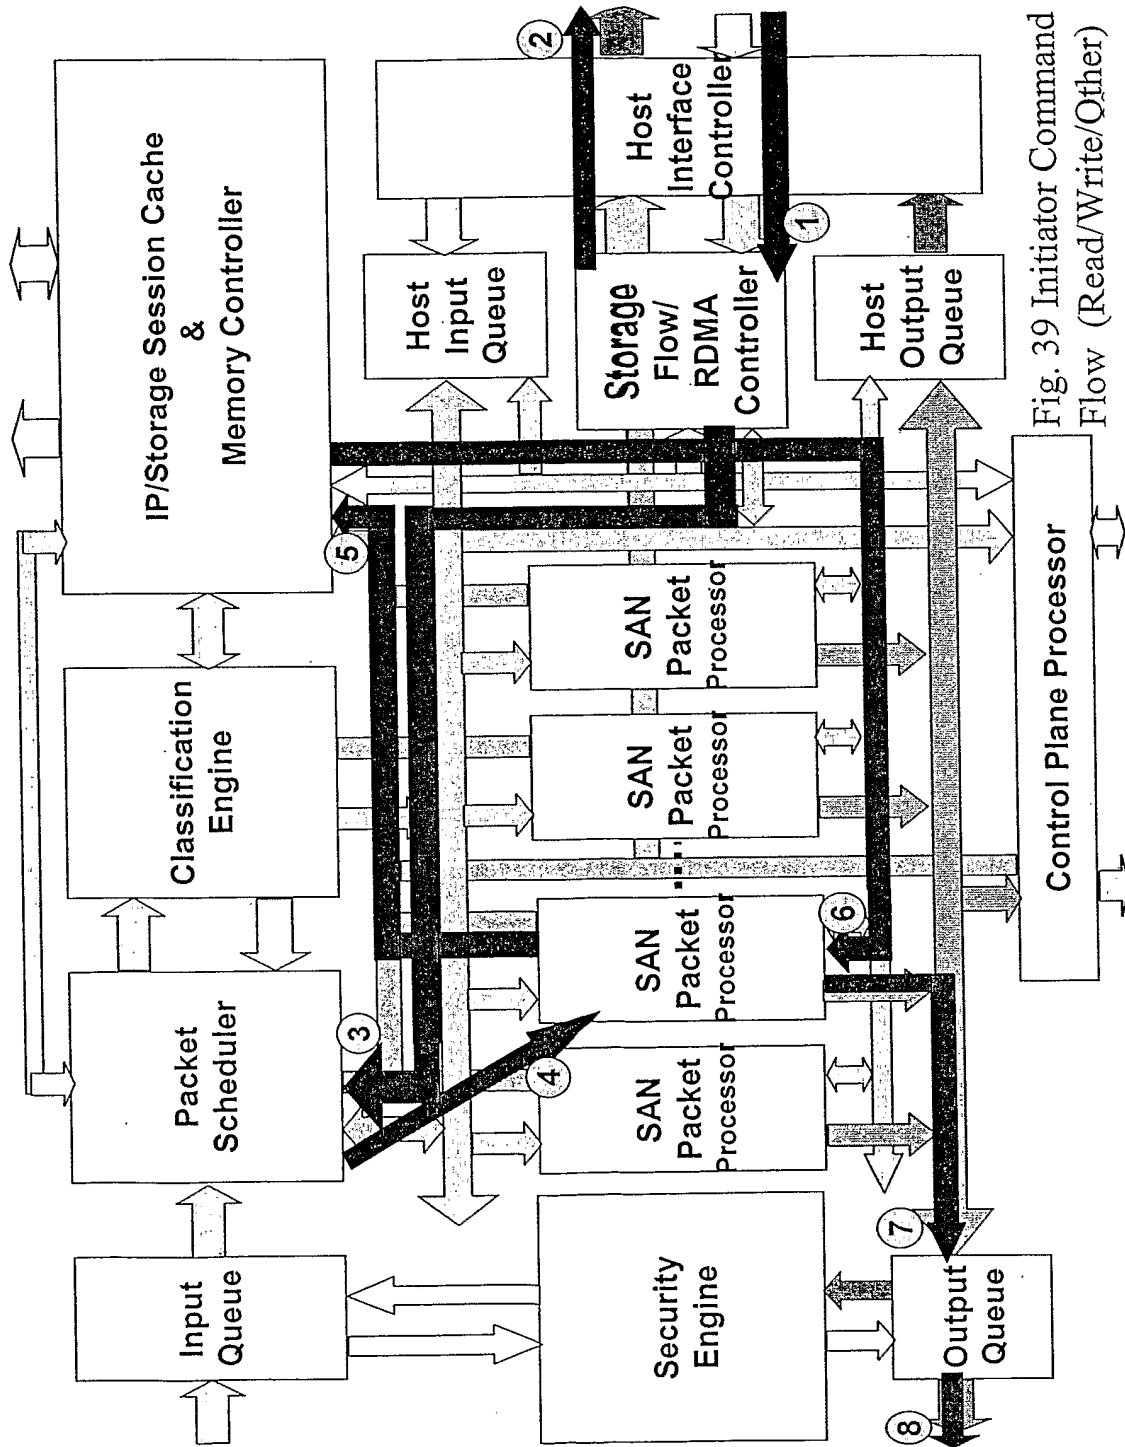


Fig. 39 Initiator Command Flow (Read/Write/Other)




Fig. 40 Read Data Flow

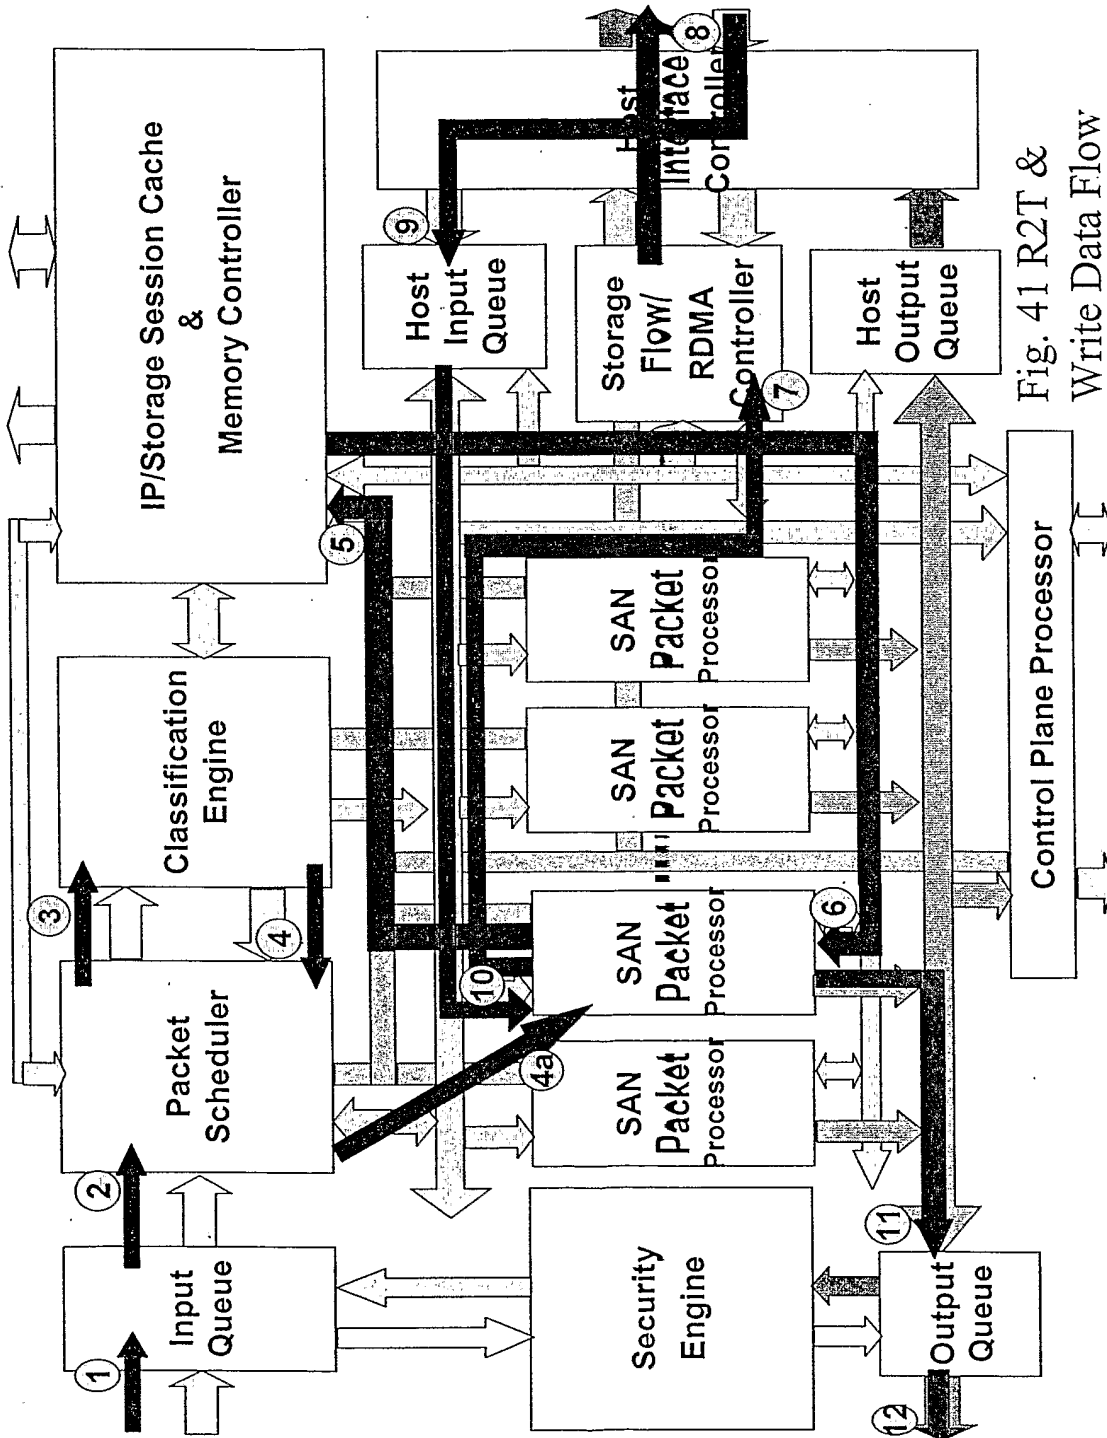


Fig. 41 R2T & Write Data Flow




Fig. 42 Read Data Flow (Secure)




Fig. 43 R2T & Write Data Flow (secure)




Fig. 44 RDMA Buffer Adv




Fig. 45 RDMA Write Flow

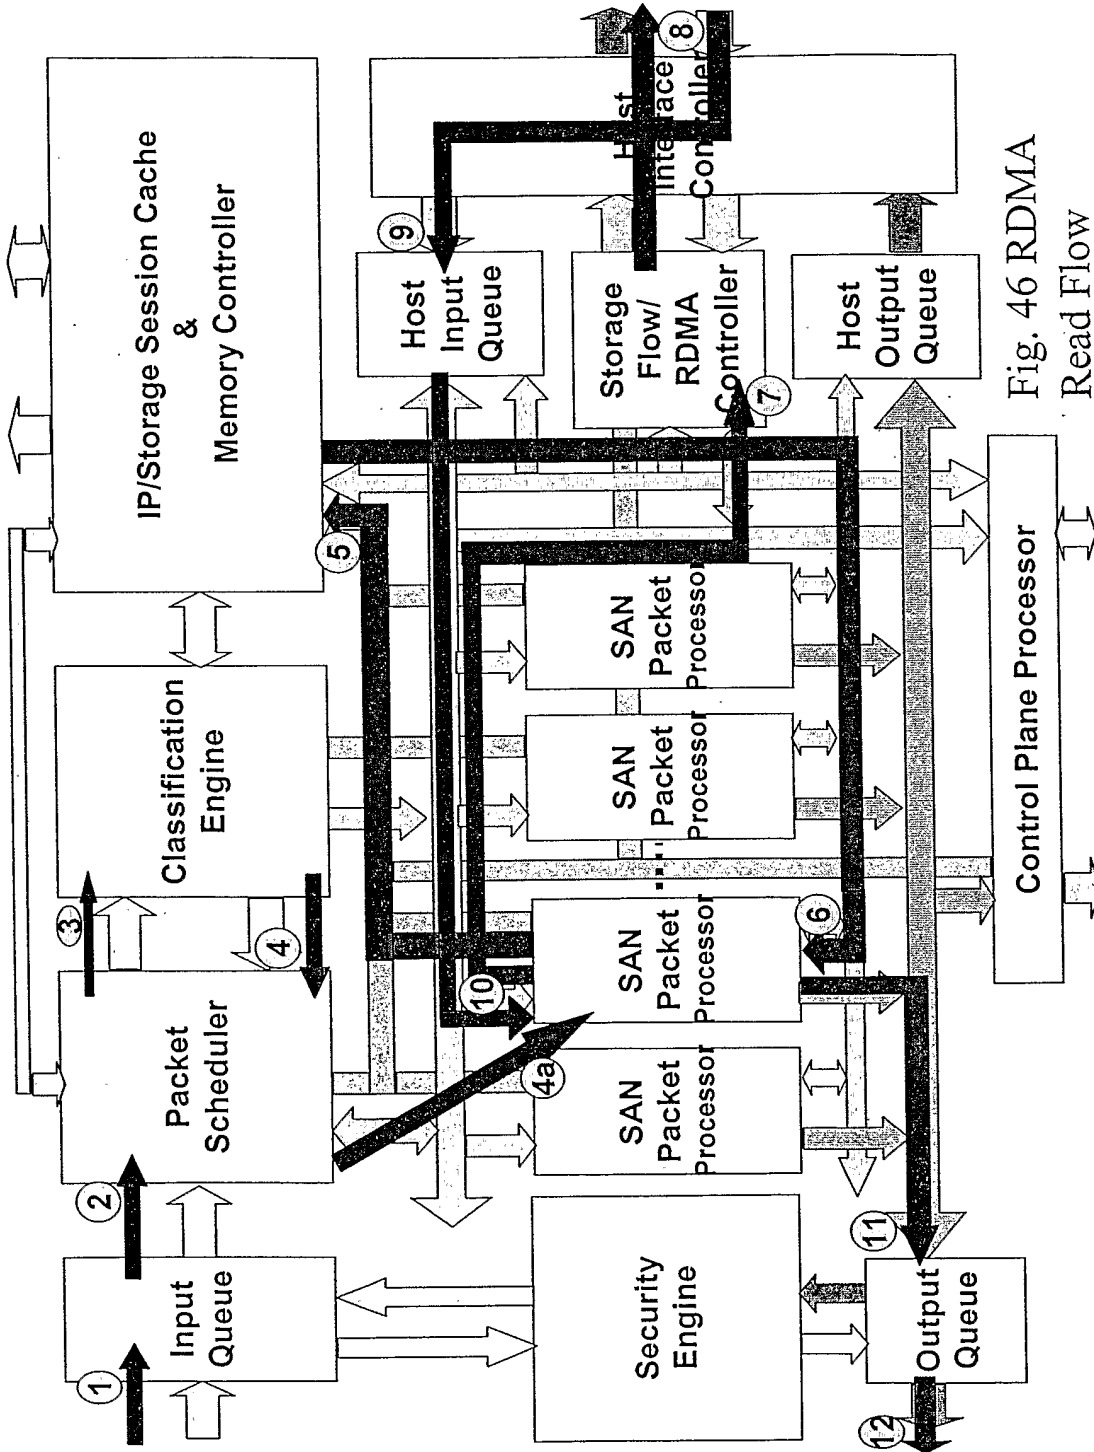


Fig. 46 RDMA Read Flow

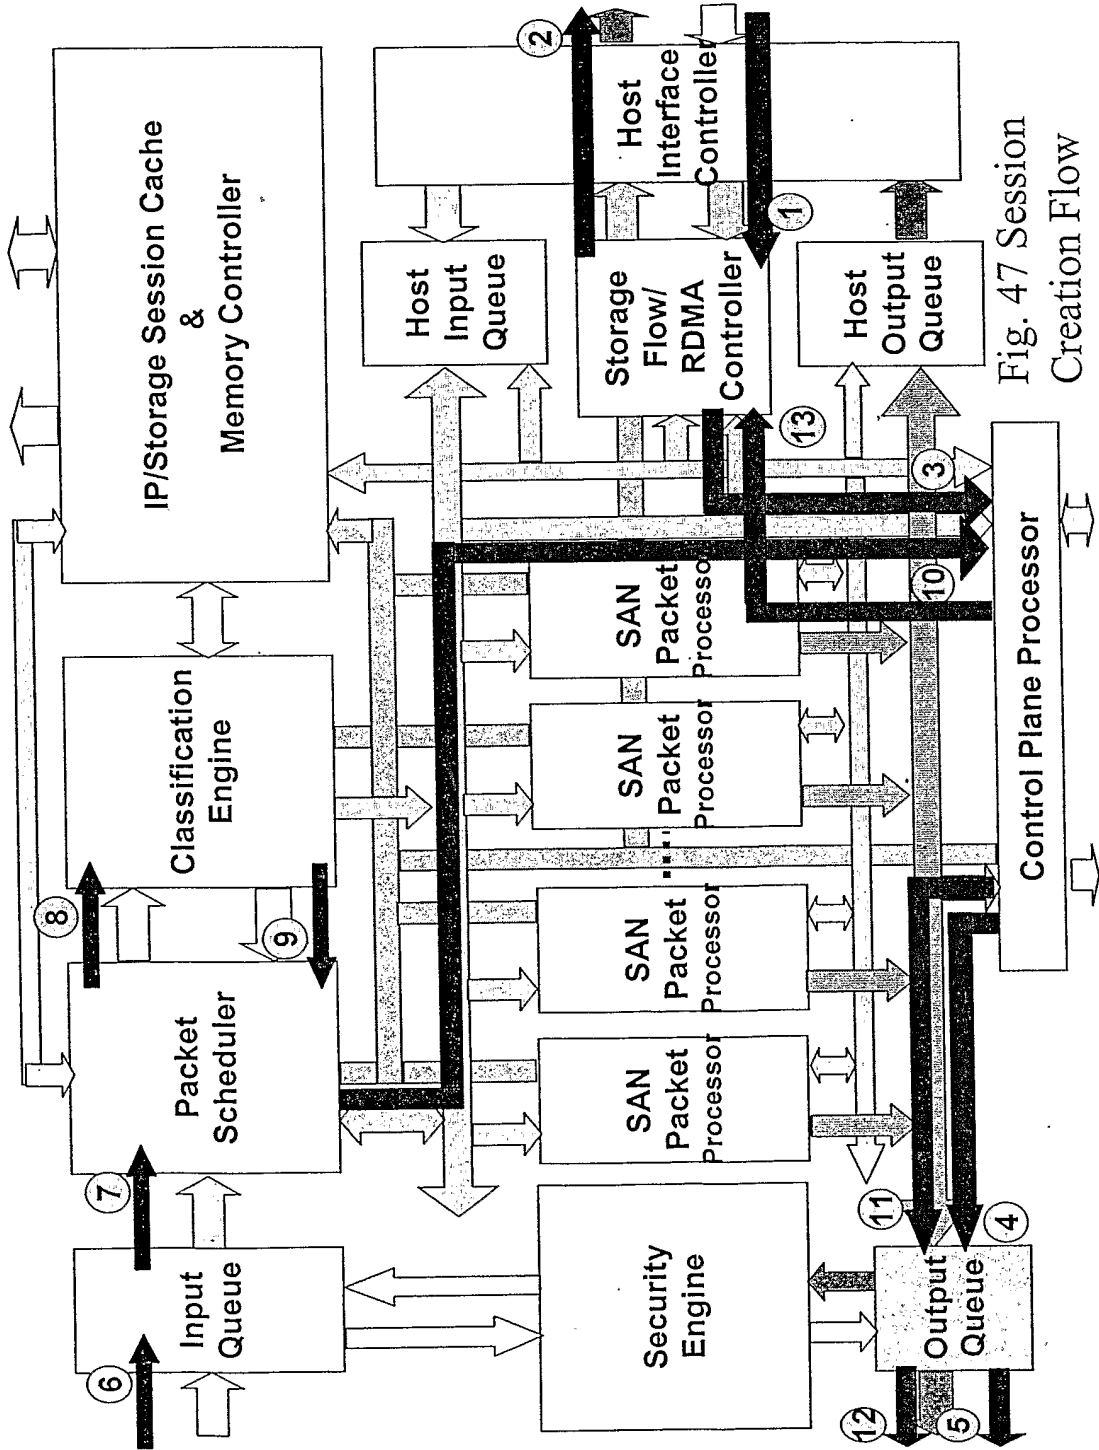


Fig. 47 Session Creation Flow

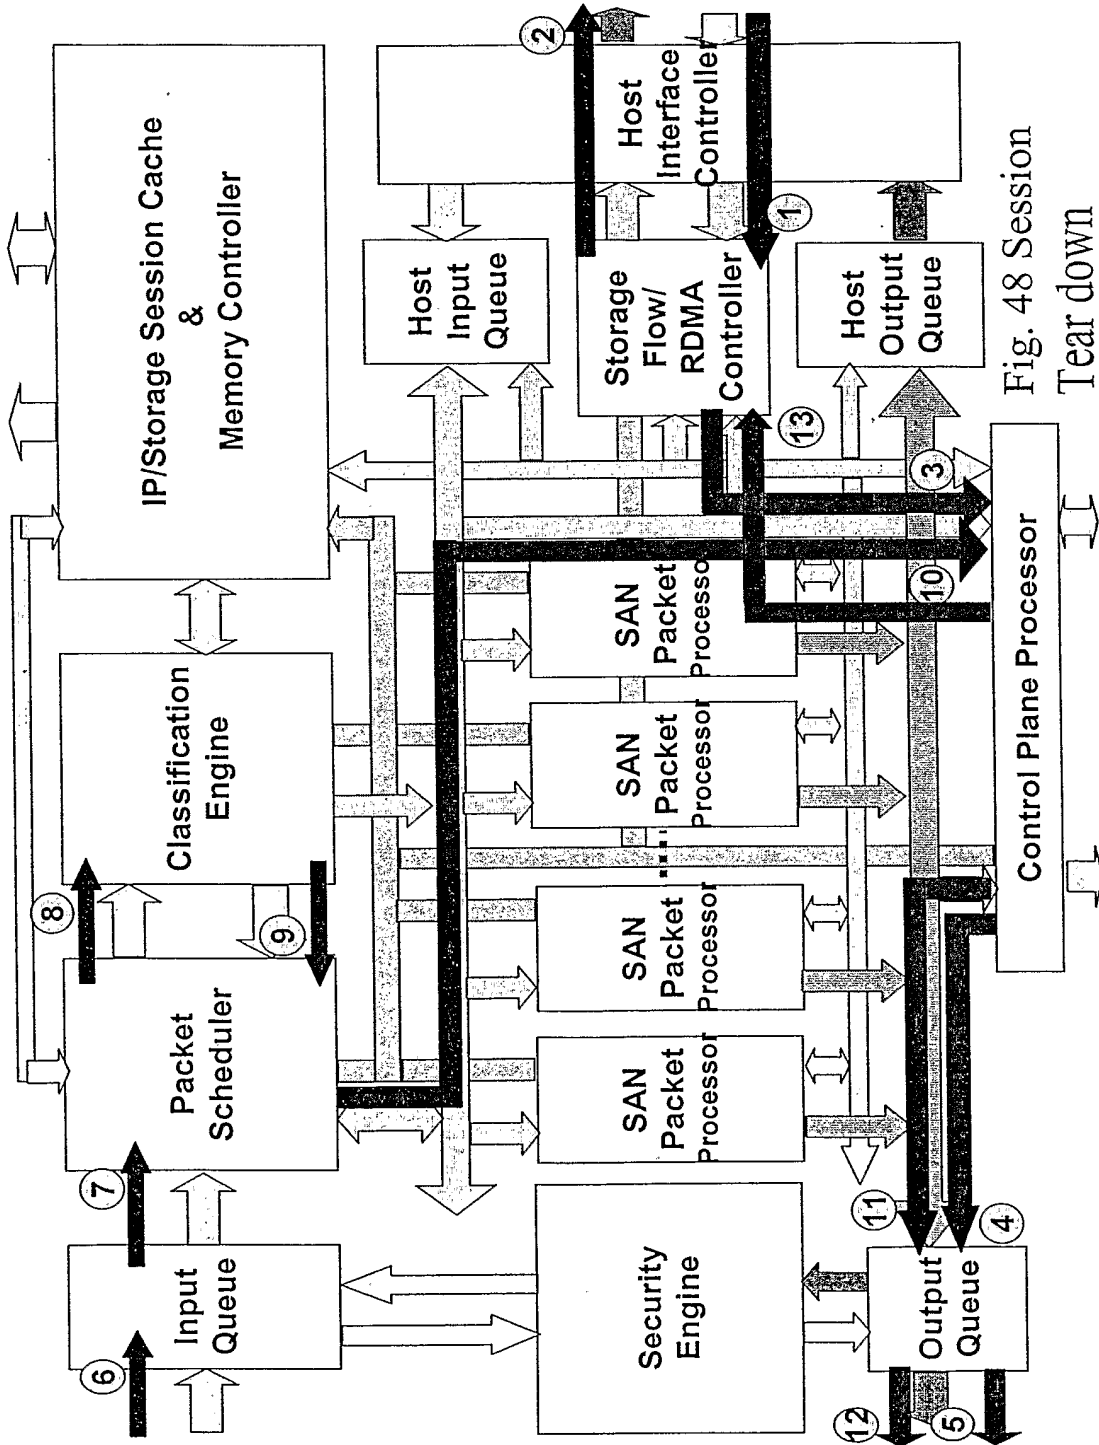


Fig. 48 Session Tear down




Fig. 49 Target Session Creation/Tear down Flow




Fig.50 Target Write Flow (Send R2T)




Fig. 51 Target Write Flow (Receive Write)




Fig. 52 Target Read Flow

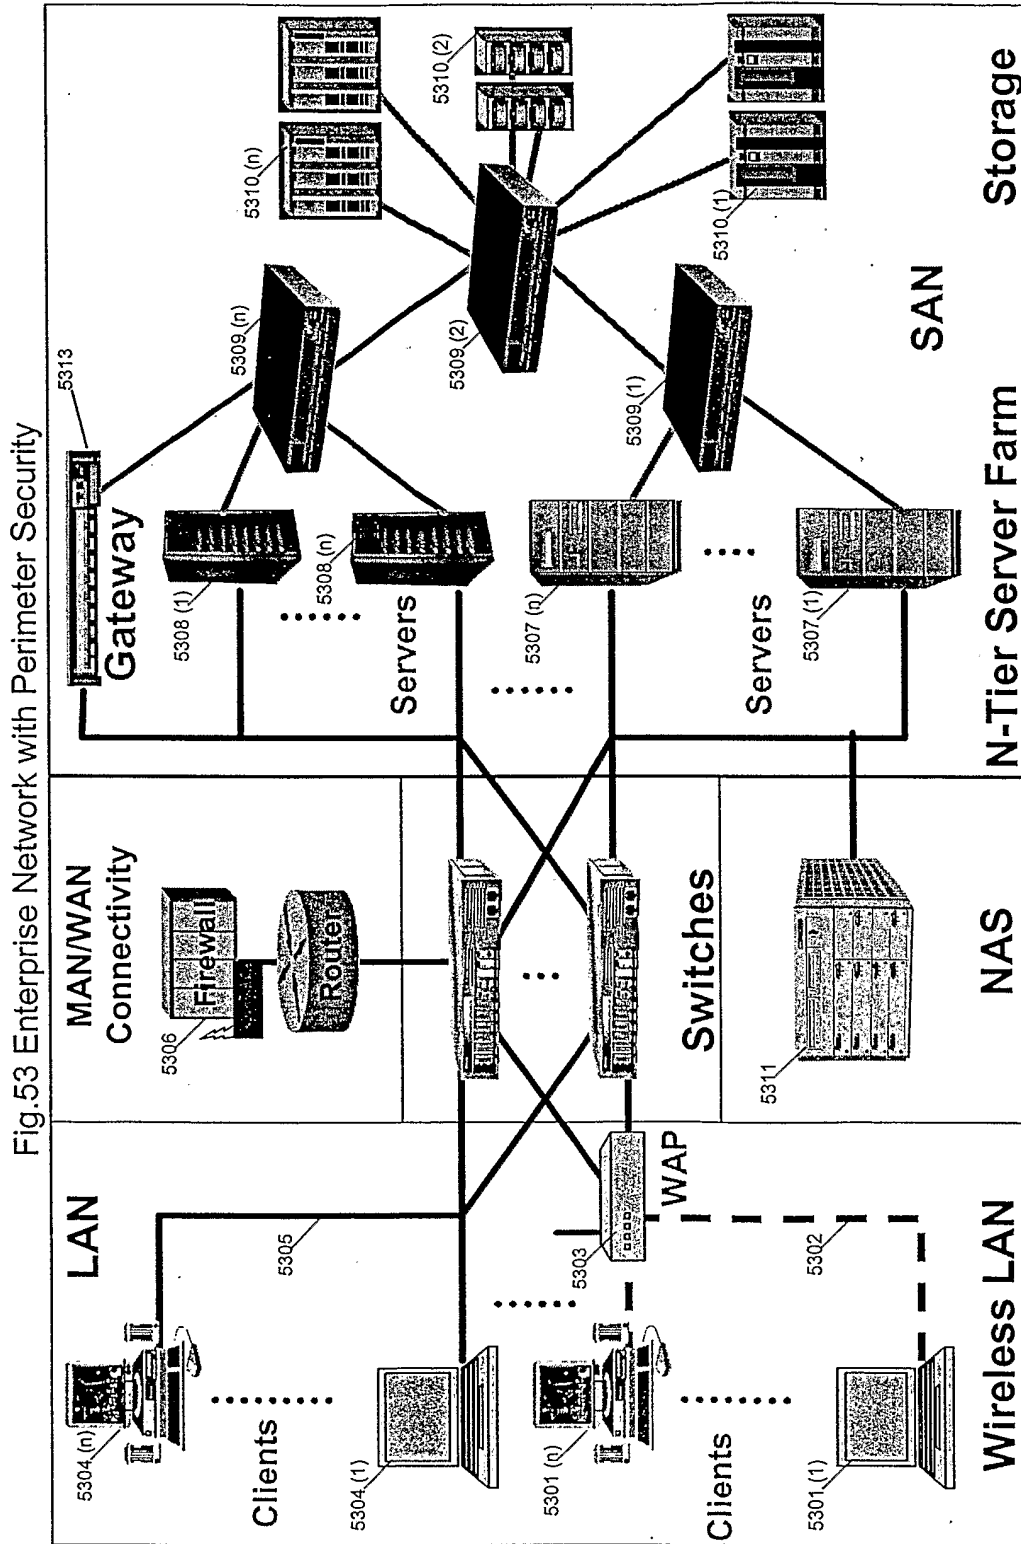


Fig.53 Enterprise Network with Perimeter Security

Fig.54 Enterprise Network with Distributed Security

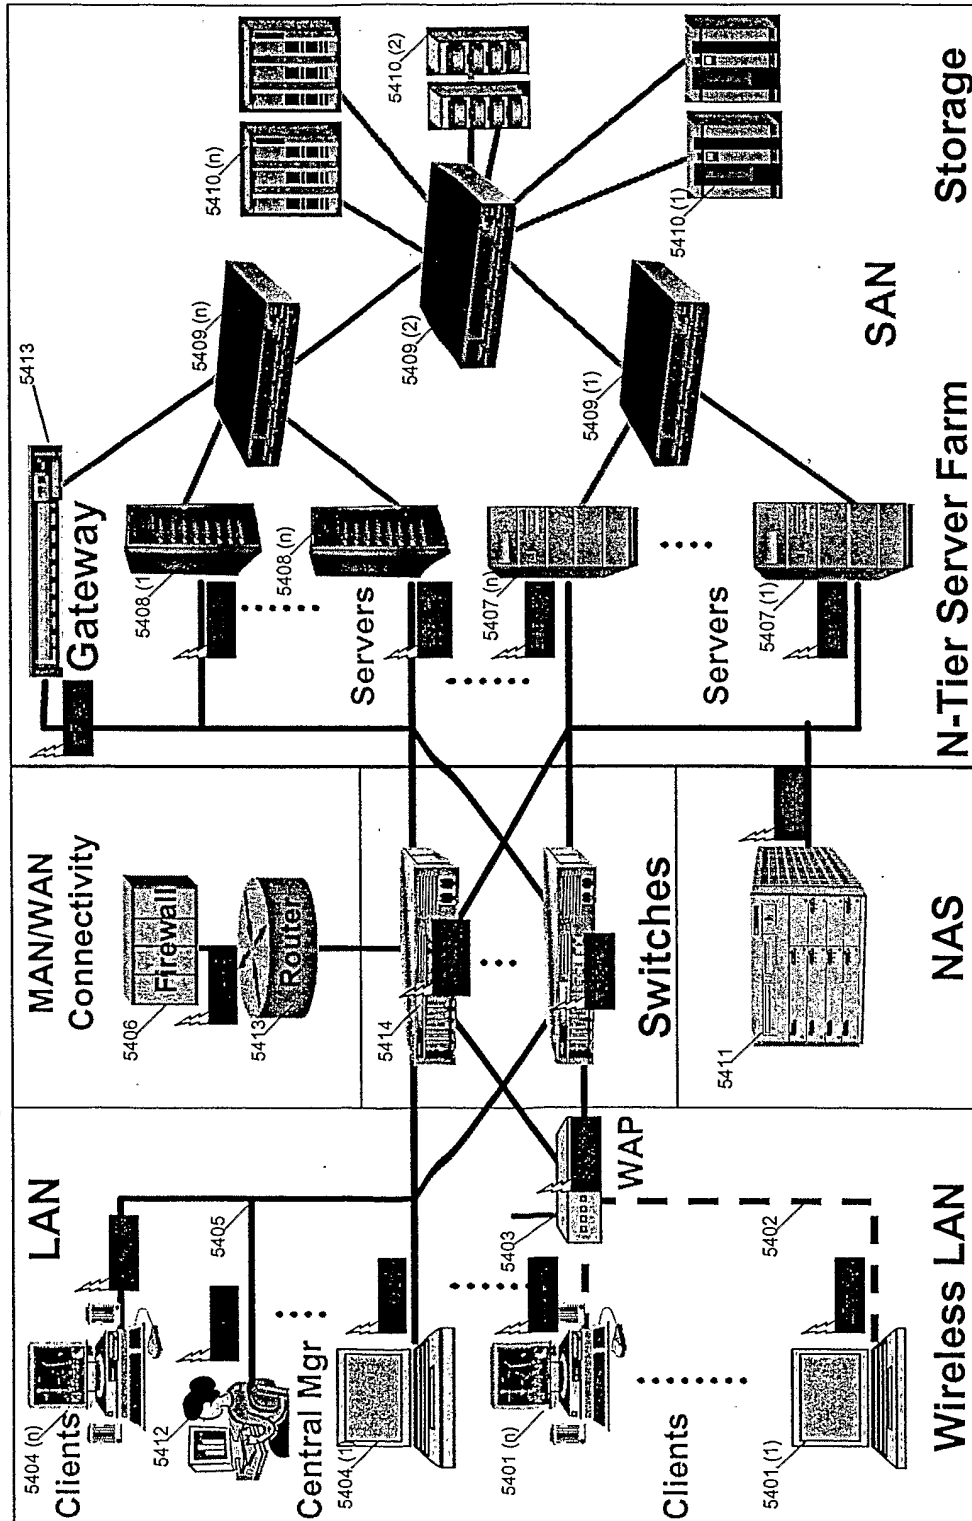



Fig.55 Enterprise Network with Distributed Security



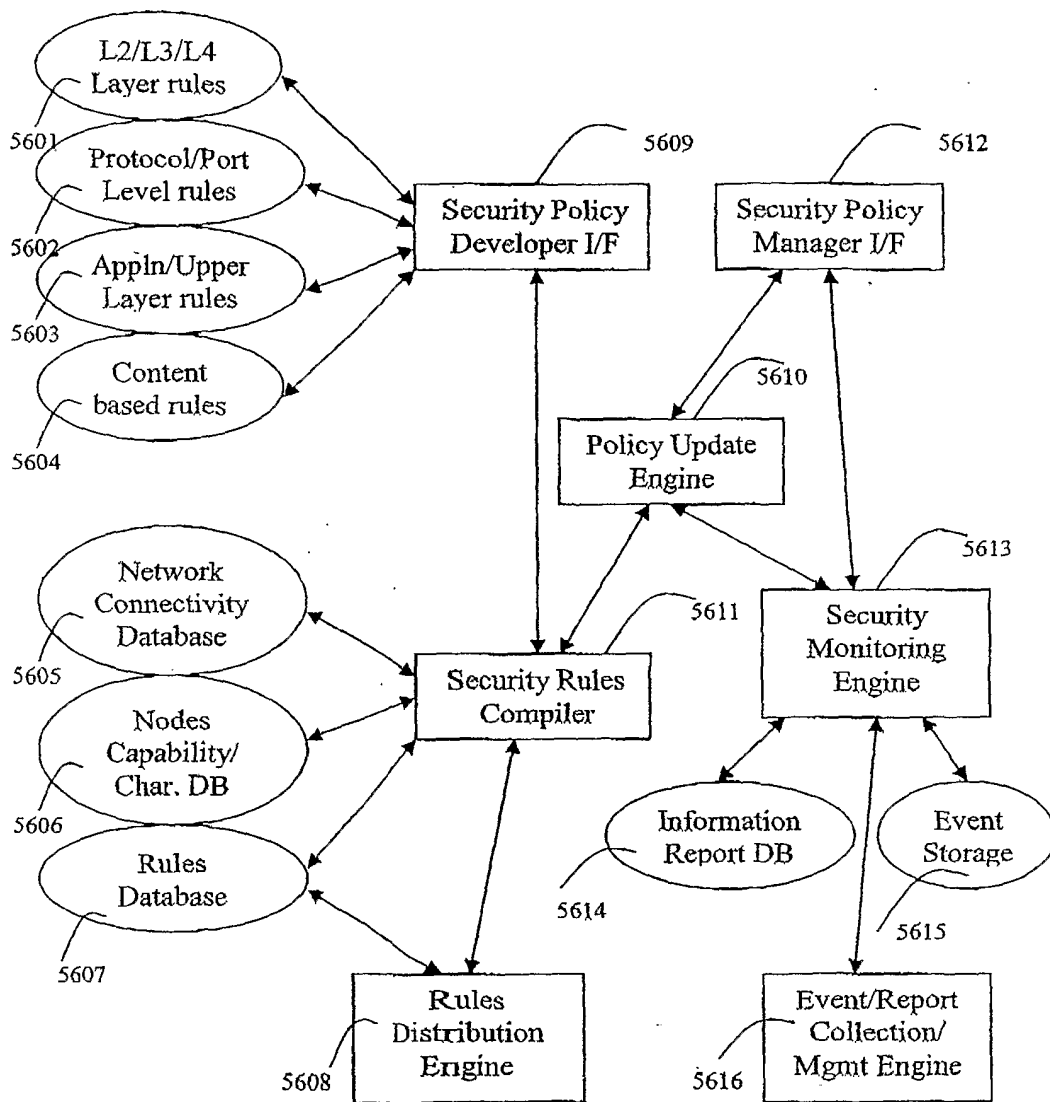


Fig. 56 Central Manager/Policy Server & Monitoring Station

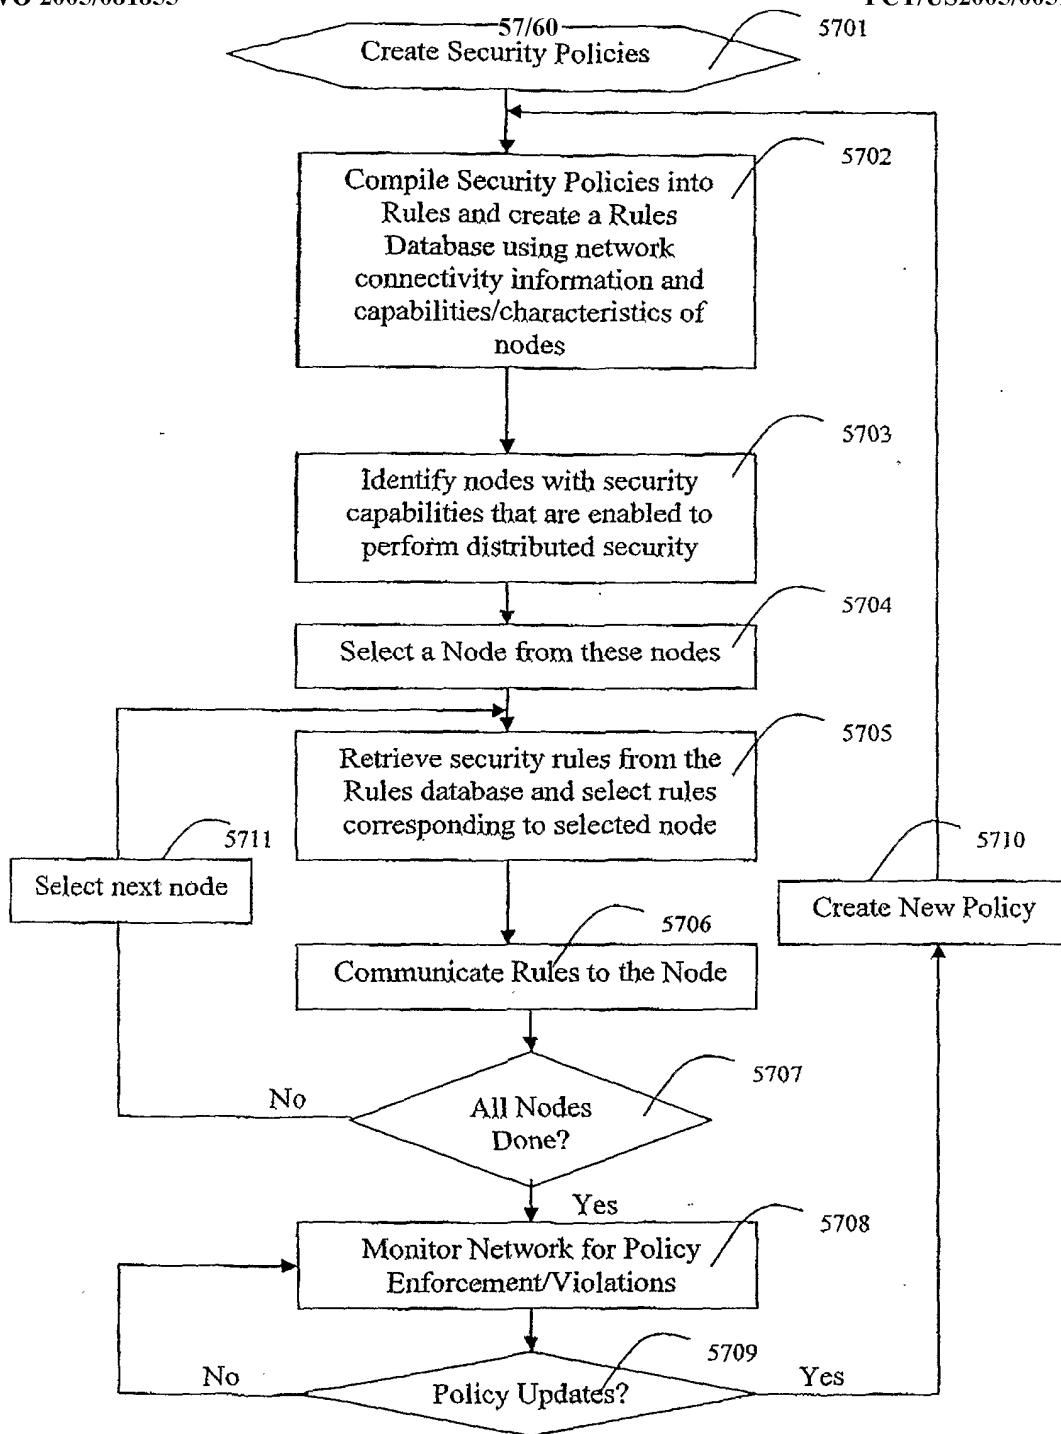


Fig. 57 Central Manager Flow

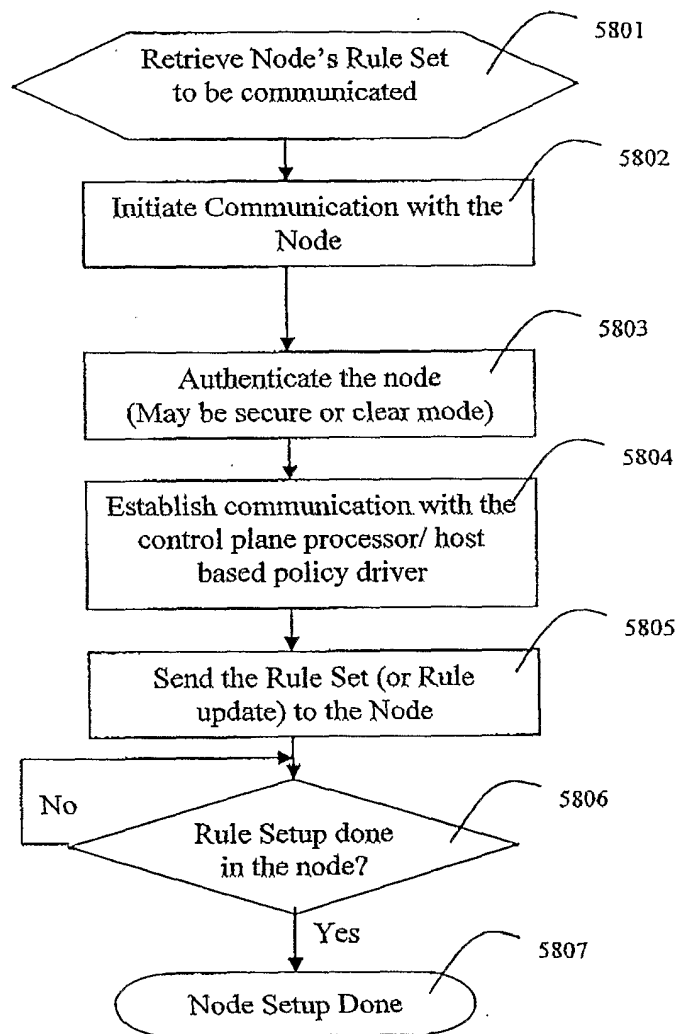


Fig. 58 Rule Distribution Flow

59/60

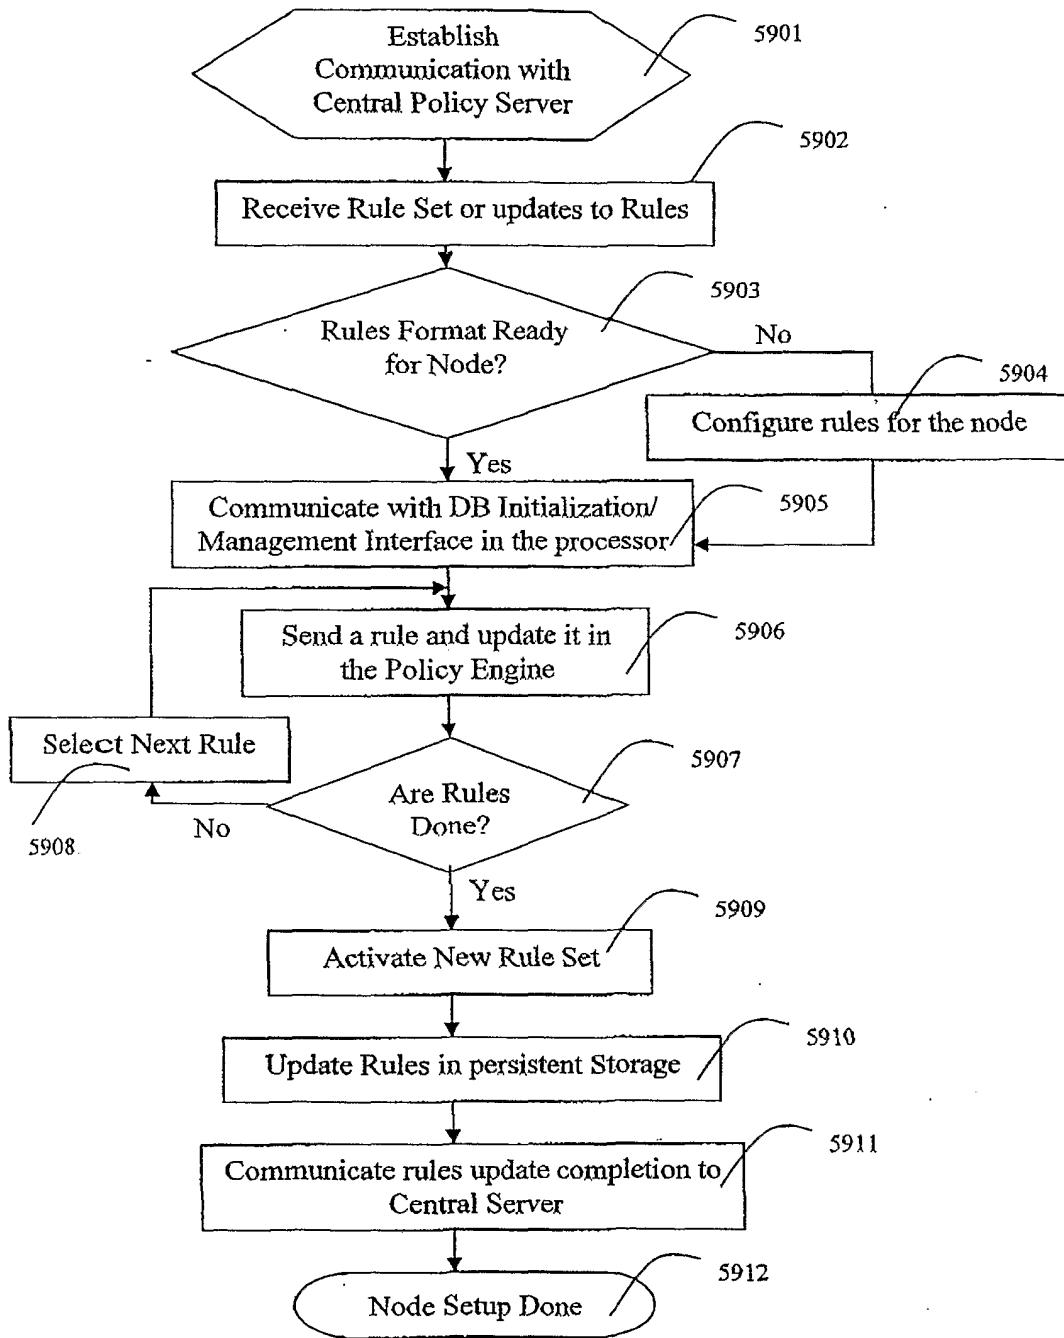


Fig. 59 Control Plane Processor/Policy Driver Flow

