

# **TOWARD CONTENT-BASED CLASSIFICATION**

## **WHITE PAPER**

**ISSUE 1: FEBRUARY 2001**

**CONTENTS**

1 PURPOSE AND SCOPE OF THIS DOCUMENT ..... 1

2 DEFINITIONS ..... 2

3 CLASSIFICATION IN NETWORKING ..... 3

4 COMPONENTS OF PACKET PROCESSING ..... 4

    4.1 FRAMING AND VALIDATION ..... 4

    4.2 PARSING..... 4

    4.3 CLASSIFICATION ..... 4

    4.4 ENCRYPTION AND DECRYPTION ..... 6

    4.5 EDITING AND SCHEDULING ..... 6

    4.6 STATISTICS COLLECTION..... 7

    4.7 FORWARDING ..... 7

5 DATA PATH FUNCTIONS ..... 8

    5.1 DATA PATH REQUIREMENTS ..... 8

    5.2 DATA PATH PROCESSING TIME ..... 10

6 SUMMARY..... 13

**LIST OF FIGURES**

FIGURE 1: APPLICATIONS ON AN 800 MIPS PROCESSOR ..... 11

FIGURE 2: PARSING AND LOOKUP OPERATION COMPLEXITY ..... 12

## **1 PURPOSE AND SCOPE OF THIS DOCUMENT**

This paper examines the importance of packet classification and content processing in networking applications.

## **2 DEFINITIONS**

ACL	Access Control List
ATM	Asynchronous Transfer Mode
CIDR	Classless Inter Domain Routing
DIP	Destination Internet Protocol
DWDM	Dense Wavelength Division Multiplexing
FDDI	Fiber Distributed Data Interface
HDLC	High-Level Data Link Control
HSSI	High-Speed Serial Interface
HTTP	Hyper-Text Transfer Protocol
L2TP	Layer Two Tunneling Protocol
IDS	Intrusion Detection and Security
IPSec	IP Security
ISDN	Integrated Service Digital Network
MAC	Media Access Control
MPLS	Multi-protocol Label Switching
NAT	Network Address Translation
POS	Packet Over Sonet
PPTP	Point-to-Point Tunneling Protocol
QoS	Quality of Service
SAP	Service Advertising Protocol
TCP	Transmission Control Protocol
TOS	Type of Service
UDP	User Datagram Protocol
URI	Universal Resource Identifier
URL	Uniform Resource Locator
VBS	Visual Basic Script
VOIP	Voice Over IP
VPN	Virtual Private Network

### **3 CLASSIFICATION IN NETWORKING**

The rapid growth of Internet usage is revolutionizing the networking industry. Practically everyone has access to a wide range of Internet services, including entertainment, special interest group content, news, secure business transactions, data storage, shared file access and information exchanges. These heterogeneous services, ranging from time-critical audio and video streaming to volume-intensive file/data transfers, rely on a variety of network transport characteristics. Differentiation of the service quality level for each type of access is extremely important as mission critical traffic co-mingles with best effort traffic. Additionally, enterprise computing and networking resources must be protected from malicious attacks as access to an increasing variety of services becomes more widespread.

Network equipment vendors are transforming multiple, special-purpose boxes into single application-convergent products that aggregate different network functions and applications to address the complexities of managing an increasingly wide range of network services. Network functions such as routing, firewall protection, network address translation (NAT), Quality of Service (QoS), and metering and traffic shaping are becoming an integral part of today's network equipment. In addition, more complex transport policies and strategies are evolving to handle information differentiation efficiently.

In the early stages of Internet development, packet-forwarding decisions were based on simple look-up operations at the Layer 2 and Layer 3 fields of a packet header. Today, multiple look-up operations using the Layer 3 and Layer 4 fields, as well as flow/session state maintenance, are commonplace features in networking equipment. For next generation network equipment, the focus of attention is on analysis and classification of packets in the Layer 5 through Layer 7 fields (packet payload). Simply identifying the packet through the header information is no longer sufficient. Network applications such as load balancing, web-switching, intrusion detection and security (IDS) require deep packet analysis. It is apparent to system implementers that examination of packet data payload is essential in order to address the needs of next generation networks.

As the traffic demand on different network systems increases, the need for high-speed network infrastructures intensifies. Recent advances in optical technology, such as DWDM and Lambda switches, are making high bandwidth pipes a reality. Lines rate of 2.5 Gbps (OC-48) and 10 Gbps (OC-192) are now available and 40 Gbps (OC-768) is within sight. In the foreseeable future, optical fiber will carry terabits of data to and from subscriber premises. The result of the increasing complexity of network traffic, combined with ever-increasing wire-speeds, impacts next generation equipment design. These new boxes must have architecture that provides intricate and extremely fast packet processing capabilities. Intelligent, high performance packet classification is key.

## **4 COMPONENTS OF PACKET PROCESSING**

Packets flowing through the network are, in essence, fragments of information exchanged between source and destination application layers. Packets can be thought of as being composed of Layer 2 through Layer 4 headers followed by data payload. Upon receiving a packet from the media, the network equipment would perform packet processing. Packet processing can include the following operations: framing and validation, parsing, classification, encryption/decryption, editing and scheduling, statistics gathering, and forwarding. The packet processing stages are described below.

### **4.1 Framing and Validation**

The bit stream is received from the network medium and converted to distinguishable frames or cells. Examples of frames/cells are Ethernet, Token Ring, FDDI, HSSI, HDLC, POS, and ATM. These data packets are also verified for validity.

### **4.2 Parsing**

Various header fields from the packet, such as MAC addresses, IP network addresses, TCP or UDP port numbers, etc. are extracted, depending on the protocol and application requirement of the network system. Parsing involves identification of the nature of the structure of the packet according to standard protocols.

### **4.3 Classification**

The basic function of network equipment is to forward an inbound packet received on one port to another port that brings the packet closer to its destination machine. The forwarding decision is based on packet classification results. The classification stage identifies the packet by inspecting various fields and segments within the header and data payload and comparing them to a statically or dynamically defined look-up database. Depending on the application requirements, classification may consist of a single look-up or it may comprise a sequence of various look-up operations. Packet classification is governed by sets of policy rules that resolve the intent of certain network functions and applications. Some typical classification functions are described below.

#### **MAC Address Forwarding and Learning**

The packets are classified according to the Layer 2 destination Media Access Control (MAC) address and, optionally, the Virtual LAN tag. The MAC address identifies the packet destination, either final or intermediate, within a LAN segment, by performing a look-up operation against a valid MAC address table. The learning process of the source MAC address for the database update is also involved in the classification task.

## IP Routing

The IP routing packet classification task generally involves a look-up operation on the Layer 3 destination IP (DIP) address field in the IP header. The DIP identifies the intended recipient of the packet. The Classless Inter Domain Routing (CIDR) implements segmentation of the Internet address into a hierarchical, logically addressable group of sub-networks. The CIDR addressing scheme requires longest-prefix-lookup operation to identify the routing path.

## Access Control List (ACL) Filtering

At a high-level, a network is a mesh of interconnected computers, workstations, and networking equipment. In a rudimentary sense, access between one network node and another is unrestrained. However, to bring in a degree of control, all network domain packets flowing through the system are provided with service, such as bandwidth, access, etc. according to the network's governing access control policy. The access control policy is implemented by filtering packets at strategic points within the network. In IP networks, ACL-based filtering involves distinguishing a packet's 5-tuple information: Layer 3 source and destination IP addresses; Layer 4 source and destination Ports; and the Layer 4 Protocol. This multi-field packet classification is a component of a firewall function that virtually encapsulates and filters in- and out-bound traffic in a network domain. Another application that embodies multi-field classification is QoS functionality. QoS defines the level of access service in terms of bandwidth, delay, jitter, etc. Network services can be differentiated according to application, subscriber discrimination, destination identification, information content and other rules. With IP, certain schemes utilize the Type of Service (TOS) field or the DiffServ tag in the option field in addition to the 5-tuple information and other pertinent header fields.

## Flow Identification and State Maintenance

Packets that share common application layer or session layer information are grouped into a flow. Classification of packets is required to identify the flow to which they belong and to maintain and track the specific state of the associated connection/session. Identification of flow membership is required to maintain the connection path and to guarantee the service level required by the application or subscriber. For firewall applications, in order to establish the qualified point of filtering, not only should inspection on multiple header fields be performed but the state of the connection must also be examined.

## Content Processing

The Internet is being propelled by countless web-based services. Web-based applications primarily employ the Hyper-Text Transfer Protocol (HTTP). The host names and web pages' Uniform Resource Locator (URL) are present within the data payload of an HTTP packet. An HTTP packet is normally identified by the destination port number

that is contained in the header. However, the web host specification and the Universal Resource Identifier (URI), which identify the path of access, must be parsed from the packet data payload. Because these fields do not reside in a specific offset within the packet, complex search operations must be performed to locate and classify them. These processes are required in applications such as web-caching, web-switching, server-sticky connections and other web-based functions. More advanced content processing involves parsing for and locating the cookie field to identify the subscriber. Recent, high profile attacks on popular websites have brought into focus the urgent need for deploying Intrusion Detection and Security (IDS) systems. IDS requires parsing of complex regular expressions involving tokens, variable-sized strings, and wild cards within the data payload.

Thus, it is evident that packet content processing requires substantial processing resources and time.

### **Multi-dimensional Classification**

Driven by service providers' needs, special-purpose networking functions and applications are converging in several network appliance implementations and equipment classes that sit on the network domain borders. With convergence of different functions, multi-dimensional classification consisting of multi-field header lookups and significant content analysis are necessary.

#### **4.4 Encryption and Decryption**

Packet classification is the groundwork for encryption and decryption. Certain networking services, such as secure e-commerce transactions and Virtual Private Networks (VPNs), require encryption and decryption of packet contents. The Secure Socket Layer (SSL) protocol provides encrypted communications on the Internet. For VPNs, based on protocols such as IPSec, L2TP and PPTP, packets flowing through the network are selectively encrypted/decrypted according to the tunneling mechanism selected. Packets need to be classified in order to identify the tunneling scheme and the associated encryption and decryption processes.

#### **4.5 Editing and Scheduling**

Packet classification is the foundation for all subsequent packet-processing stages. The results of the packet analysis and classification become the basis for packet editing. Editing may involve field modification, as in NAT, or label/tag content manipulation for packet forwarding schemes. Field modification also requires regeneration of the correct redundancy code. Packet tagging may be required to administer traffic scheduling within the system's inter-connection fabric or when higher priority packets need to be dispatched over lower priority packets.

## **4.6 Statistics Collection**

Statistics collection is a significant element in network management and control. It is essential for monitoring network resource usage, performance characteristics, traffic analysis, and service level analysis. Statistics are gathered from the results of various stages of the packet classification process.

## **4.7 Forwarding**

When application-driven packet classification is complete, forwarding decisions are executed. In a network data plane, packets are forwarded via a switch fabric to the qualified port at the network equipment interface.

## **5 DATA PATH FUNCTIONS**

Based on the packet-forwarding model described above, it is clear that the classification stage carries the bulk of the processing burden. As multiple transport policies and forwarding functions have emerged and evolved, the intricacy of the packet classification task has been compounded. Multi-dimensional classification involving multi-field header lookups and data payload parsing is now required in various network functions. Sophisticated classification is the vital element in achieving high-speed networks.

### **5.1 Data Path Requirements**

Traditionally, the routing function only involved a longest-prefix-match lookup operation on the IP Layer 3 destination address based on the CIDR scheme. In large-scale IP networks, the Multi-Protocol Label Switching (MPLS) protocol was defined to combine the Layer 2 switching and Layer 3 routing functions. This is similar to the strategy employed in ATM or Frame Relay switching. The key concept in label switching is identifying and marking IP datagrams with labels and forwarding them to a modified switch or router that uses the labels to switch the datagrams through the network. This data path function normally requires large lookup tables.

Some variants of routing consist of examining the type of service (TOS) field and the DiffServ label field, which identify the TOS required for the connection, in conjunction with the QoS function. The primary goals of QoS include:

- control of the resources including equipment and facilities, bandwidth, jitter and latency,
- improved loss characteristics, and
- efficient differentiation between mission-critical applications from best-effort accesses.

Network infrastructures, like the ISDN, integrate data, audio and video under one communication channel. Recent protocol definitions that integrate the various media services are Voice Over IP (VOIP), which simultaneously carries voice and data traffic, and the H.323 standard for sending audio and video on the Internet and within intranets. In terms of service delivery, these real-time applications and other interactive Internet services are time-critical. Although there is no requirement for large lookup tables for QoS, sophisticated sequential parsing and classification on multiple fields is required.

A higher level of QoS implementation requires identification of flows of the application traffic and services. Certain network accesses can be considered as conversation rather than pure connection. Identifying and isolating traffic flow in the midst of inter-mixed information exchanges is a vital step in implementing efficient management of resources. Flow management of applications like Oracle, SAP, Informix, NetMeeting and others,

requires extensive flow identification, tracking and maintenance. In FTP access, well-known or assigned port numbers are generally used during connection establishment; however, port numbers would then be dynamically re-assigned when the information exchange phase is entered. Hence, dynamic update of policy rules is necessary as flows are created or deleted and as the state transitions. Edge equipment normally maintains a database of a million or more flows and flow association requires a lookup operation against this large database.

In order to protect network resources from malicious users or accidental catastrophic accesses and safeguard confidential information, firewalling must be employed. The firewall function is achieved with the help of packet filtering along with per session/flow state maintenance. The basic filtering is controlled by a defined policy that is laid down in an ACL. Typically, classification for packet filtering involves examining the 5-tuple fields in a packet. A filtering policy rule requires that each field of the 5-tuple be compared with arithmetic and relational operators such as bit masking, equality, range compare, etc. Firewall policies also implement active filtering based on the state of a session.

IDS function extends the firewall implementation. Service providers, enterprises, institutions and campuses are paying more attention to intrusion detection and network system security. Recent events show that networks can be vulnerable to malicious attacks, either intentional or accidental. An attack can be launched using several methods, including brute force resource manipulation or latent programs that unleash attacks by accessing unsuspecting internal authorized users. One form of attack takes advantage of the weaknesses of application layer software. Examples are manipulating Visual Basic Scripts (VBS) for macro programs, and exploiting mailing systems and server page protocol idiosyncrasies. For IDS applications, scrutiny of the packet data payload is executed by parsing some defined regular expression strings. Regular expressions can be comprised of multiple tokens, case insensitive short and long patterns, wild cards, variable white spaces and string chains that exist at no particular location within the data payload. It is evident that IDS requires tremendously complex packet classification tasks.

Analysis of packet data payload content is also widely employed in several network appliances directly associated with web servers. In World-Wide-Web-based applications, network transactions are actually specified beyond the Layer 4 header. Since only the assigned port numbers, for example port 80 for HTTP, are normally used, in-depth parsing of the packet payload is required. The URL contained in the data payload of a packet specifies the location of a web page object on the Internet. In such applications, the web page is specified by a Hypertext Markup Language (HTML) document, and the URL string forms one hyperlink target. Information contained in a URL can be:

- the transport protocol, such as HTTP, FTP, gopher, Wide Area Information Server (WAIS),
- the hostname specification,

- the Universal Resource Identifier (URI), and
- the query of fragment identifier.

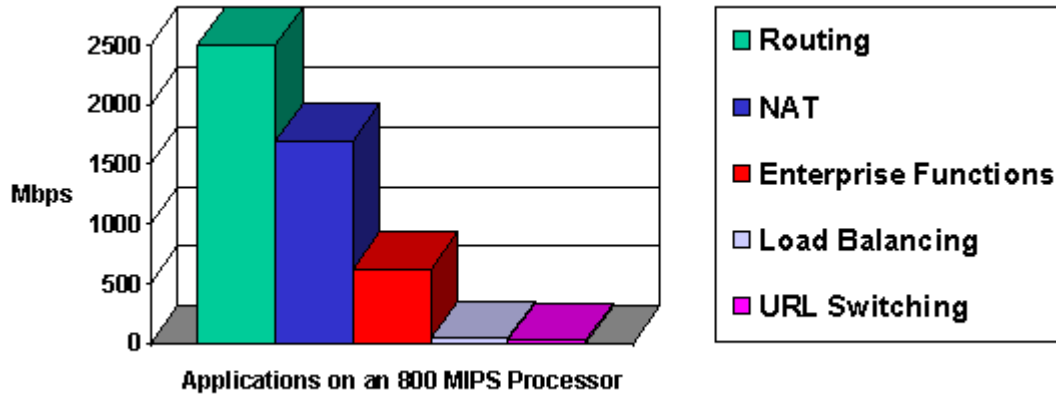
The URL string must be parsed and classified to properly implement server redirection and alignment, as in web-switching and web-caching. Not only is the classification task complex but a large lookup table is normally required. Another data payload component in web-based applications is the HTTP cookie. A cookie is a packet of information, exchanged between an HTTP server and a web browser, that can contain any arbitrary information used to maintain the state between otherwise stateless HTTP transactions. The packets must be classified based on the cookie contents to maintain sticky connections when, for example, an e-commerce or e-shopping client must be connected to a single and unique server throughout the transaction.

With the growth of public backbone infrastructures, the Internet is becoming the transport channel in all remote networking connections. Many enterprises and organizations are out-sourcing network function services to Metro application service providers. Enterprise functions like firewall, NAT, and QoS are now being provided and maintained by more centralized providers. The complexity of the packet classification on the aggregation points is magnified by multiple functions and heavy traffic load. In addition, VPNs are being widely deployed for secure end-to-end private connections over the Internet. Certain encryption strategies and tunneling protocols, such as IPSec, L2TP and PPTP, are employed on selected packets coming to and from VPN subscribers. In conjunction with these advances, deployment of server farms is ceaselessly ramping. Load balancing, which requires multi-stage classification operations, is necessary for implementing traffic engineering and congestion control.

## **5.2 Data Path Processing Time**

An analysis of various scenarios for implementing different networking functions provides an insight into data path processing time utilization. The figure below shows the throughput achieved when networking functions such as routing, Network Address Translation (NAT), enterprise functions (including firewall, QoS, etc.), load balancing, and URL switching are implemented. The figures presented have been estimated for a system based on an 800 MIPS processor without any special hardware assist for classification.

**Figure 1: Applications on an 800 MIPS Processor**

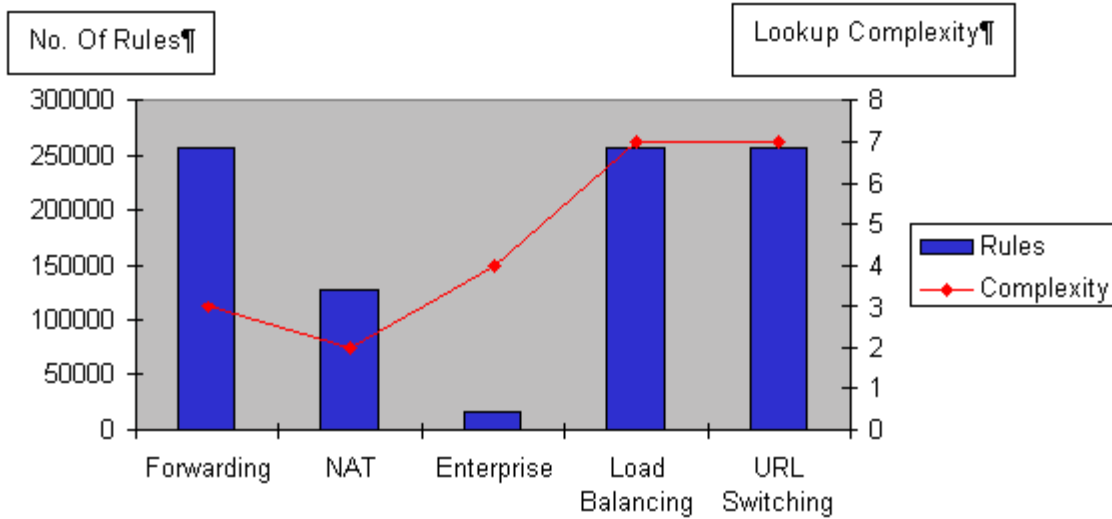


As evident in the graph, processing throughput performance drops as the complexity of classification increases. Packet classification complexity is largely influenced by the following factors:

- number of lookups required,
- the volume of the lookup entries, and
- the complexity of parsing and lookup operations.

Parsing and lookup operation complexity magnifies when deep-packet data payload content analysis is implemented. Figure 2 illustrates this.

**Figure 2: Parsing and Lookup Operation Complexity**



It is clear that achieving wire-speed performance in network equipment necessitates sophisticated classification hardware assist. Platform acceleration focusing on the classification task is the prime ingredient for achieving wire-speed performance in a system.

## **6 SUMMARY**

The networking industry is facing great challenges and complex issues in deploying high-speed networks. What is hampering significant progress is not only the creation and development of faster media but also the realization of high-speed packet processing systems. In the packet-processing model, the bottleneck is the classification task. A specialized solution must be employed to off-load the classification task from the processor. Existence of a packet classifier in a data plane allows the processor to perform the tasks required before and after classification. These tasks include checking packet validity and packet parsing before the classification is performed, and packet content manipulation and forwarding based on classification results. Further, to provide maximum benefit, the packet classification has to be done, not only on the basis of the headers, but also the contents of each packet.

**CONTACTING PMC-SIERRA, INC.**

PMC-Sierra, Inc.  
105-8555 Baxter Place Burnaby, BC  
Canada V5A 4V7

Tel: (604) 415-6000

Fax: (604) 415-6200

Document Information: [document@pmc-sierra.com](mailto:document@pmc-sierra.com)

Corporate Information: [info@pmc-sierra.com](mailto:info@pmc-sierra.com)

Application Information: [apps@pmc-sierra.com](mailto:apps@pmc-sierra.com)

(604) 415-4533

Web Site: <http://www.pmc-sierra.com>

None of the information contained in this document constitutes an express or implied warranty by PMC-Sierra, Inc. as to the sufficiency, fitness or suitability for a particular purpose of any such information or the fitness, or suitability for a particular purpose, merchantability, performance, compatibility with other parts or systems, of any of the products of PMC-Sierra, Inc., or any portion thereof, referred to in this document. PMC-Sierra, Inc. expressly disclaims all representations and warranties of any kind regarding the contents or use of the information, including, but not limited to, express and implied warranties of accuracy, completeness, merchantability, fitness for a particular use, or non-infringement.

In no event will PMC-Sierra, Inc. be liable for any direct, indirect, special, incidental or consequential damages, including, but not limited to, lost profits, lost business or lost data resulting from any use of or reliance upon the information, whether or not PMC-Sierra, Inc. has been advised of the possibility of such damage.

© 2001 PMC-Sierra, Inc. PMC-Sierra™ is a trademark of PMC-Sierra, Inc.

PMC-2002233

Issue date: February 2001

---