US 8,165,024 B2

Use of DPI to extract and forward application characteristics

Exemplary of:
WAN and Application Optimization Solutions

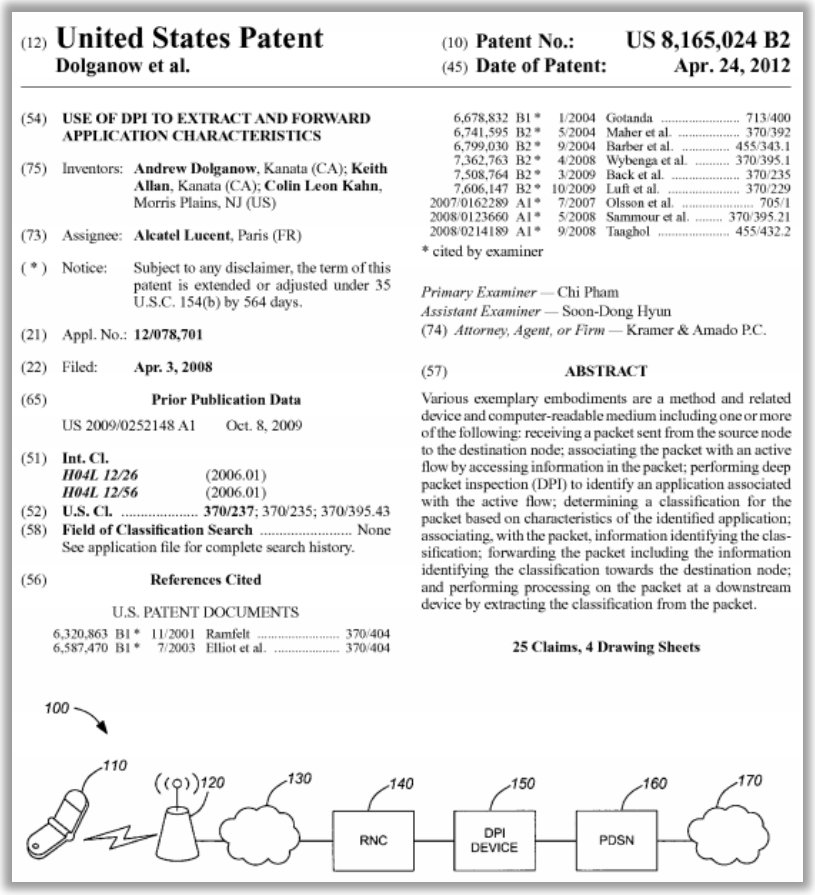

(12) **United States Patent**
Dolganow et al.

(10) **Patent No.: US 8,165,024 B2**
(45) **Date of Patent: Apr. 24, 2012**

(54) **USE OF DPI TO EXTRACT AND FORWARD APPLICATION CHARACTERISTICS**

(75) Inventors: **Andrew Dolganow**, Kanata (CA); **Keith Allan**, Kanata (CA); **Colin Leon Kahn**, Morris Plains, NJ (US)

(73) Assignee: **Alcatel Lucent**, Paris (FR)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 564 days.

(21) Appl. No.: **12/078,701**

(22) Filed: **Apr. 3, 2008**

(65) **Prior Publication Data**

US 2009/0252148 A1 Oct. 8, 2009

(51) **Int. Cl.**
H04L 12/26 (2006.01)
H04L 12/56 (2006.01)

(52) **U.S. Cl.** **370/237**; 370/235; 370/395.43

(58) **Field of Classification Search** ..... None
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

6,320,863 B1 * 11/2001 Ramfelt 370/404
6,587,470 B1 * 7/2003 Elliot et al. 370/404
6,678,832 B1 * 1/2004 Gotanda 713/400
6,741,595 B2 * 5/2004 Maher et al. 370/392
6,799,030 B2 * 9/2004 Barber et al. 455/343.1
7,362,763 B2 * 4/2008 Wybenga et al. 370/395.1
7,508,764 B2 * 3/2009 Back et al. 370/235
7,606,147 B2 * 10/2009 Luft et al. 370/229
2007/0162289 A1 * 7/2007 Olsson et al. 705/1
2008/0123660 A1 * 5/2008 Sammour et al. 370/395.21
2008/0214189 A1 * 9/2008 Taaghol 455/432.2

* cited by examiner

Primary Examiner — Chi Pham
Assistant Examiner — Soon-Dong Hyun
(74) *Attorney, Agent, or Firm* — Kramer & Amado P.C.

(57) **ABSTRACT**

Various exemplary embodiments are a method and related device and computer-readable medium including one or more of the following: receiving a packet sent from the source node to the destination node; associating the packet with an active flow by accessing information in the packet; performing deep packet inspection (DPI) to identify an application associated with the active flow; determining a classification for the packet based on characteristics of the identified application; associating, with the packet, information identifying the classification; forwarding the packet including the information identifying the classification towards the destination node; and performing processing on the packet at a downstream device by extracting the classification from the packet.

25 Claims, 4 Drawing Sheets

What is claimed is:

1. A method of processing packets sent from a source node to a destination node, the method comprising:

receiving a packet sent from the source node to the destination node;

associating the packet with an active flow by accessing information in the packet;

performing deep packet inspection (DPI) to identify an application associated with the active flow by analyzing at least one other packet;

determining a classification for the packet based on characteristics of the identified application;

inserting information identifying the classification into the packet;

forwarding the packet, including the information identifying the classification, towards the destination node such that a downstream device is enabled to perform processing of the packet by extracting the classification from the packet.

https://patents.google.com/patent/US8165024

US 8,165,024 B2 Claim 1

Exemplary of:
WAN and Application Optimization Solutions

1. A method of processing packets sent from a source node to a destination node, the method comprising:

receiving a packet sent from the source node to the destination node;

associating the packet with an active flow by accessing information in the packet;

performing deep packet inspection (DPI) to identify an application associated with the active flow by analyzing at least one other packet;

determining a classification for the packet based on characteristics of the identified application;

inserting information identifying the classification into the packet;

forwarding the packet, including the information identifying the classification, towards the destination node such that a downstream device is enabled to perform processing of the packet by extracting the classification from the packet.

IOS instrumentation provides a good starting point for creating a network performance baseline through the following components:

- NetFlow
- IPSLA
- NBAR
- CBQoS MIB

NetFlow provides a good source of traffic flow information for capturing normal and abnormal behaviors on the network. Additionally, standardized SNMP MIBs from individual devices provide basic information about the network such as traffic volume by byte, errors, utilization on interfaces, etc. NBAR, a traffic identification and classification engine built into IOS, can discover the types of applications that are present on the network. Together, NetFlow, MIBs, and NBAR provide a comprehensive baseline about the physical network and the paths application flows take as they utilize the network.

Creating response time baseline is important to the success of an IT organization in establishing service quality levels. Active and passive response time measurements are two methodologies for measuring application response times

There is no one single source of information for baselining your network and applications. IT organizations will need to use different monitoring instrumentation data in order to gain a solid understanding of the normal behavior of the applications, the network, and IT resources.

Comment: A high-level view of exemplary product's components used in network performance optimization, including Netflow and NBAR.

US 8,165,024 B2 Claim 1

1. A method of processing packets sent from a source node to a destination node, the method comprising:

receiving a packet sent from the source node to the destination node;

associating the packet with an active flow by accessing information in the packet;

performing deep packet inspection (DPI) to identify an application associated with the active flow by analyzing at least one other packet;

determining a classification for the packet based on characteristics of the identified application;

inserting information identifying the classification into the packet;

forwarding the packet, including the information identifying the classification, towards the destination node such that a downstream device is enabled to perform processing of the packet by extracting the classification from the packet.

NetFlow identifies packet flows for IP Packets by looking at a number of fields in the data packet. A flow is defined as a set of packets having common properties. NetFlow defines a flow as the combination of the following seven key-fields, which determine how a flow is identified:

1. Source IP Address
2. Destination IP Address
3. Source port number
4. Destination port number
5. Layer 3 protocol type (e.g., ICMP, TCP, UDP)
6. ToS byte
7. Logical input interface (ifIndex)

Each flow record is created by grouping packets with the same characteristics into a flow. This method of determining a flow is ideal because a large amount of network information is condensed into a database of NetFlow information called the NetFlow cache. If any of these fields are different from another flow, it is considered a different flow.

NetFlow operates by creating a NetFlow cache entry that contains information for each active flow., as illustrated in Figure 4-3.

Figure 4-3 NetFlow Cache Entry

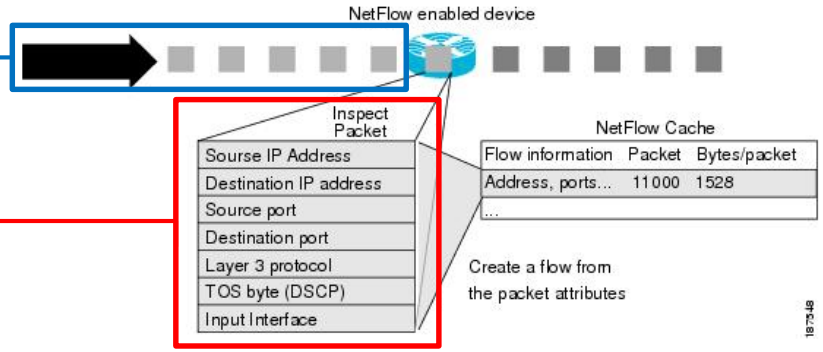

Comment: NetFlow associates the packet with a flow.

US 8,165,024 B2 Claim 1

Exemplary of:
WAN and Application Optimization Solutions

1. A method of processing packets sent from a source node to a destination node, the method comprising:

receiving a packet sent from the source node to the destination node;

associating the packet with an active flow by accessing information in the packet;

performing deep packet inspection (DPI) to identify an application associated with the active flow by analyzing at least one other packet;

determining a classification for the packet based on characteristics of the identified application;

inserting information identifying the classification into the packet;

forwarding the packet, including the information identifying the classification, towards the destination node such that a downstream device is enabled to perform processing of the packet by extracting the classification from the packet.

4.3.3 NBAR

Network Based Application Recognition (NBAR) provides network traffic classification. NBAR can recognize a very wide variety of applications by doing IP packet inspection up to OSI Layer 7. It can, for instance, differentiate between Web-based HTTP and Skype traffic, which can both use TCP port 80.

When an application is recognized, NBAR classifies the traffic for performance and accounting purposes. This function gives an operator the ability to invoke any range of services for that specific application, whether offering more or less bandwidth, latency queuing, or completely blocking certain packets.

NBAR also provides a special Protocol Discovery (PD) feature that determines which applications and protocols are traversing the network at any given time. PD captures key statistics that are associated with each protocol based on IP flows. Like NetFlow, NBAR defines IP flows as a unidirectional flow of IP packets that share the following five values:

- Source IP address
- Destination IP address
- Source port
- Destination port
- L3 protocol type

NetFlow and NBAR both leverage L3 and L4 header information. However, unlike NetFlow, NBAR also examines data from L3-L7. NBAR uses L3 and L4 and packet inspection for classification, and supports stateful inspection of dynamic-port traffic. NBAR also requires a set number of packets before making a protocol distinction.

Comment: NBAR further identifies the application by doing IP packet inspection (DPI) up to OSI Layer 7, and classifies the traffic for performance and accounting purposes. Stateful analysis requires analysis of multiple packets.

US 8,165,024 B2 Claim 1

1. A method of processing packets sent from a source node to a destination node, the method comprising:

receiving a packet sent from the source node to the destination node;

associating the packet with an active flow by accessing information in the packet;

performing deep packet inspection (DPI) to identify an application associated with the active flow by analyzing at least one other packet;

determining a classification for the packet based on characteristics of the identified application;

inserting information identifying the classification into the packet;

forwarding the packet, including the information identifying the classification, towards the destination node such that a downstream device is enabled to perform processing of the packet by extracting the classification from the packet.

5.3.2 DPI Engines

DPI engines can be co-resident in the software or can be dedicated hardware. Both have advantages and disadvantages. While a dedicated hardware provides speed and versatility, the cost of deploying such a box restricts their usage to high traffic volume environments like a Data Center or a large Enterprise Branch office.

Service Control Engine (SCE) is a good example of a dedicated hardware DPI appliance. Software-based DPI engines are cost effective, but they do consume CPU cycles and hence can be deployed only at low or medium traffic volume environments such as those found in a small or medium Enterprise Branch Office.

5.3.2.1 Service Control Engine (SCE)

Service Control Engine (SCE) is a DPI device that can do DPI and detect traffic patterns at line rates. SCE incorporates many DPI technologies such as protocol/state analysis, pattern analysis, and behavioral and heuristic analysis. SCE can also do subscriber-level classification.

SCE can be deployed in-band or out-of-band. It is typically deployed in the Data Center. If it is deployed in band, all the traffic in the network passes through SCE. If it is deployed out of band then a copy of all the traffic is passed onto SCE by the DC switch. It should be noted that in out of band mode, the SCE can only perform monitoring.

5.3.2.2 Network Based Application Recognition (NBAR)

NBAR is an application-aware classification feature in IOS. NBAR can look deep inside a packet and do stateful analysis of the information in the packet. It can recognize a number of applications, including ones that use ephemeral ports. Even with a given protocol, NBAR can look so deep inside the packets that it can categorize packets that are of the same protocol, but with different protocol-specific parameters. For example, NBAR can classify based on the URL for HTTP packets and based on ICA traffic for CITRIX ICA.

Typically, QoS and NBAR are used in conjunction. NBAR is used to recognize specific applications and QoS is used to mark them and provide appropriate treatment based on the markings.

Comment: NBAR is a DPI engine.

US 8,165,024 B2 Claim 1

1. A method of processing packets sent from a source node to a destination node, the method comprising:

receiving a packet sent from the source node to the destination node;

associating the packet with an active flow by accessing information in the packet;

performing deep packet inspection (DPI) to identify an application associated with the active flow by analyzing at least one other packet;

determining a classification for the packet based on characteristics of the identified application;

inserting information identifying the classification into the packet;

forwarding the packet, including the information identifying the classification, towards the destination node such that a downstream device is enabled to perform processing of the packet by extracting the classification from the packet.

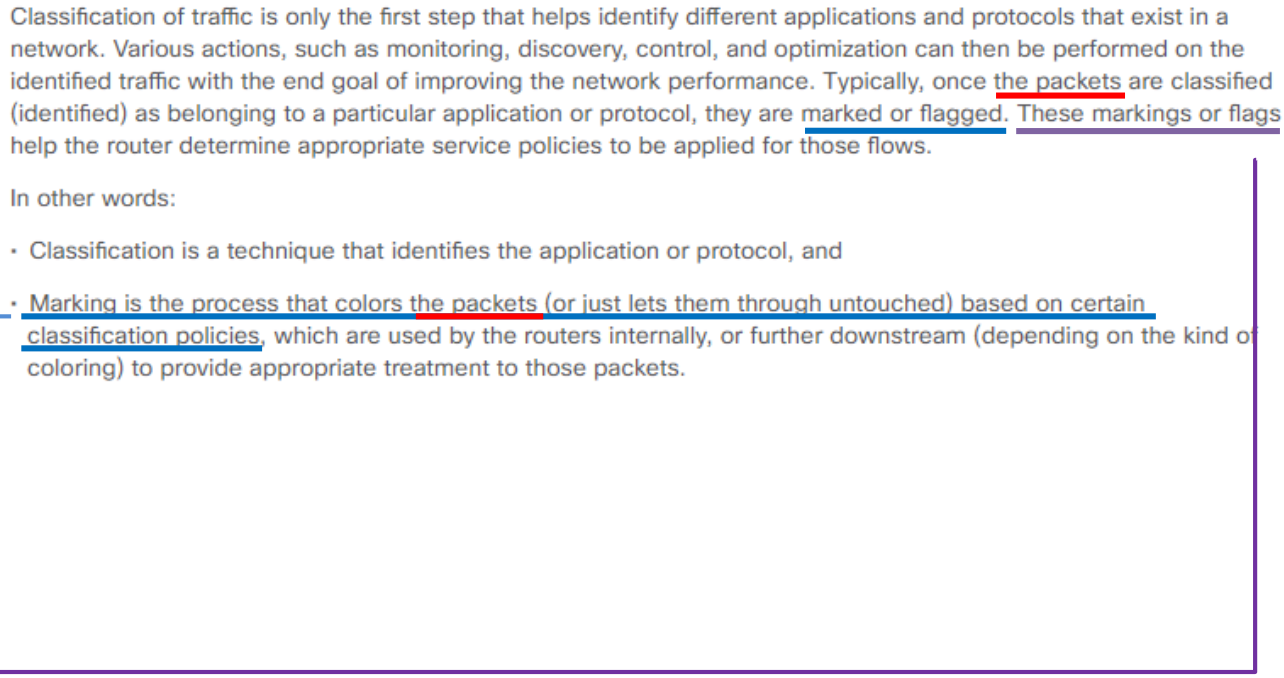

Classification of traffic is only the first step that helps identify different applications and protocols that exist in a network. Various actions, such as monitoring, discovery, control, and optimization can then be performed on the identified traffic with the end goal of improving the network performance. Typically, once the packets are classified (identified) as belonging to a particular application or protocol, they are marked or flagged. These markings or flags help the router determine appropriate service policies to be applied for those flows.

In other words:

- Classification is a technique that identifies the application or protocol, and
- Marking is the process that colors the packets (or just lets them through untouched) based on certain classification policies, which are used by the routers internally, or further downstream (depending on the kind of coloring) to provide appropriate treatment to those packets.

Comment: Once a packet is classified, it is marked (e.g. DSCP packet marking inserted into its IP header as the classification information) before forwarding towards its destination.

US 8,165,024 B2 Claim 1

1. A method of processing packets sent from a source node to a destination node, the method comprising:

receiving a packet sent from the source node to the destination node;

associating the packet with an active flow by accessing information in the packet;

performing deep packet inspection (DPI) to identify an application associated with the active flow by analyzing at least one other packet;

determining a classification for the packet based on characteristics of the identified application;

inserting information identifying the classification into the packet;

forwarding the packet, including the information identifying the classification, towards the destination node such that a downstream device is enabled to perform processing of the packet by extracting the classification from the packet.

5.4 Packet Markings

Once the flow and packets have been identified, they need to be marked so that appropriate service policies can be applied on them. The markings or flags can be set in a number of ways: for IP, Type of Service (ToS) or Differentiated Services Code Point (DSCP); for Ethernet packets, VLAN priority, etc. However, L3 markings are the most widely used method.

5.4.2 L3 Packet Markings

Similar to Layer 2 headers, the IP header has fields that can be used to classify traffic into treatment groups. The most widely used L3 marking techniques are Type of Service (ToS) and DSCP. Figure 5-5 shows a typical IP header.

Figure 5-5 \IP Header

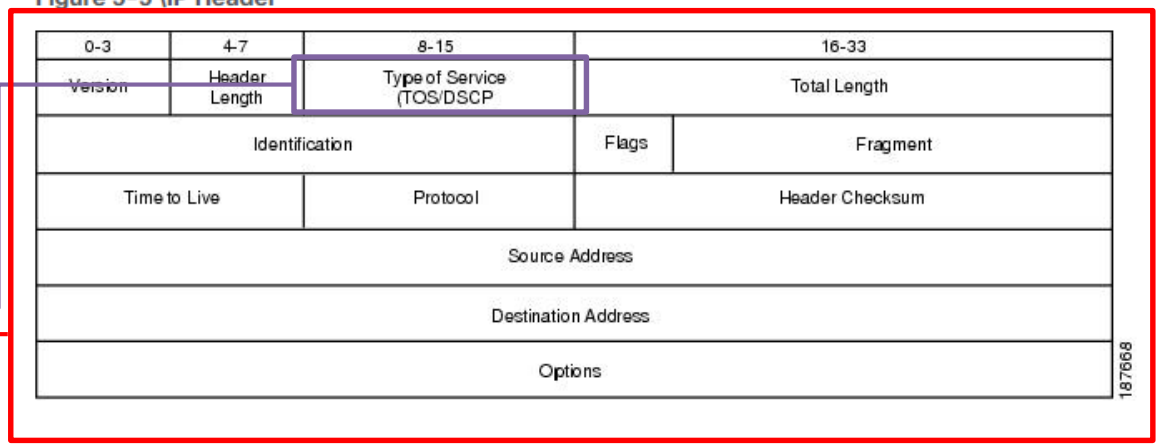

Comment: Figure 5-5 shows the header of the IP packet

US 8,165,024 B2 Claim 1

1. A method of processing packets sent from a source node to a destination node, the method comprising:

receiving a packet sent from the source node to the destination node;

associating the packet with an active flow by accessing information in the packet;

performing deep packet inspection (DPI) to identify an application associated with the active flow by analyzing at least one other packet;

determining a classification for the packet based on characteristics of the identified application;

inserting information identifying the classification into the packet;

forwarding the packet, including the information identifying the classification, towards the destination node such that a downstream device is enabled to perform processing of the packet by extracting the classification from the packet.

Classification of traffic is only the first step that helps identify different applications and protocols that exist in a network. Various actions, such as monitoring, discovery, control, and optimization can then be performed on the identified traffic with the end goal of improving the network performance. Typically, once the packets are classified (identified) as belonging to a particular application or protocol, they are marked or flagged. These markings or flags help the router determine appropriate service policies to be applied for those flows.

In other words:

- Classification is a technique that identifies the application or protocol, and
- Marking is the process that colors the packets (or just lets them through untouched) based on certain classification policies, which are used by the routers internally, or further downstream (depending on the kind of coloring) to provide appropriate treatment to those packets.

Comment: A downstream device can apply appropriate treatment, e.g. WAN optimization or application acceleration, based on the marking.

US 8,165,024 B2 Claim 1

Exemplary of:
WAN and Application Optimization Solutions

1. A method of processing packets sent from a source node to a destination node, the method comprising:

receiving a packet sent from the source node to the destination node;

associating the packet with an active flow by accessing information in the packet;

performing deep packet inspection (DPI) to identify an application associated with the active flow by analyzing at least one other packet;

determining a classification for the packet based on characteristics of the identified application;

inserting information identifying the classification into the packet;

forwarding the packet, including the information identifying the classification, towards the destination node such that a downstream device is enabled to perform processing of the packet by extracting the classification from the packet.

Figure 3-2 End-to-End WAN and Application Optimization

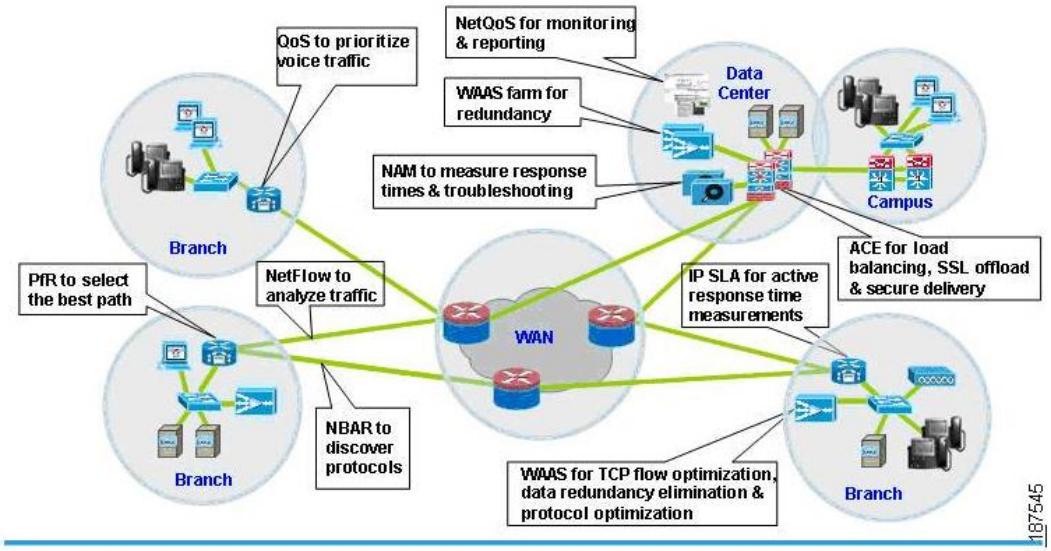

As discussed in the preceding sections, WAN and application optimization is not a single technique. It is a collection of techniques and tools working cooperatively to improve application performance. For example, in Figure 3-2 various techniques and tools are enabled in different places in the network.

Inside the branch, NetFlow and NBAR are enabled in the branch access router to provide extensive visibility into the network and applications. With visibility into the applications and their utilization, IT operations can apply QoS policies in the branch router to establish transmission priorities of the application mix. A WAAS appliance can be deployed to apply a suite of WAN optimization and application acceleration technologies to dramatically improve application performance. When the branch has dual links, performance can be further enhanced by selecting the optimal path by using PfR.

Preservation of Source TCP/IP Information

Many optimization products create tunnels through routers and other networking devices, which result in a loss of source TCP/IP information in the optimized data. This loss of TCP/IP information often disrupts important network services (such as QoS and NBAR), and can disrupt proper operation of traffic analysis tools such as NetFlow and security products and features such as ACLs and IP-based firewalls.

Unlike other optimization products, WAAS seamlessly integrates into your network and preserves all TCP/IP header information in the traffic that it optimizes, so that your existing analysis tools and security products are not compromised.

Comment: Once a packet is marked (NBAR, QoS) in the IP header, this information is preserved and used by a downstream WAAS appliance for, e.g. WAN optimization, application acceleration, protocol optimization etc.