Application is More than Header Deep

Frank Yue, 2012-02-11

In the world of application awareness and providing advanced services based on application context, it has become increasingly difficult to segregate network traffic based on the application or protocol being used. Application Delivery Controllers (ADC) are designed to manage the delivery of content through availability, performance and content optimization. Now this is starting to sound like a marketing pitch, but bear with me.

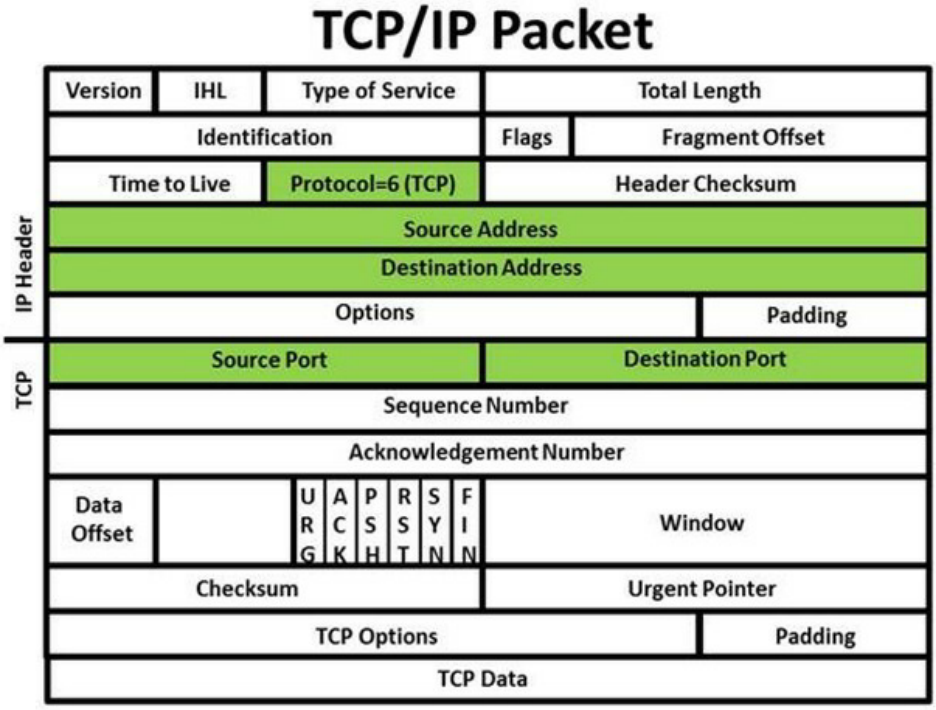

Originally, ADCs would inspect and manipulate content delivery through the inspection of packet headers. They identified the nature of the content request based on a 5-tuple. A 5-tuple typically consists of:

- Source IP address
- Destination IP address
- Source port (TCP/UDP)
- Destination port (TCP/UDP)
- IP Protocol

These five parameters can generally define a unique connection between two hosts. Applications were usually determined by the destination port. When the Internet was created, ports were mapped to common (and uncommon) application protocols based on what was known at the time. When traffic as seen with a destination of TCP port 80, it was well assumed that this traffic was HTTP. Originally, this list was managed through the regular publication of an RFC, the latest, RFC 1700, was published in 1994. Today, this list is managed by the Internet Assigned Numbers Authority, IANA, online.

The problem is that application developers and malicious users have found that content is just content and many applications are just a method for delivering that content. We have found botnets using ICMP and DNS for sharing information. Instead of sending legitimate application traffic, they are padding the payload of these packets with their subversive content.

I am willing to go out on a limb and claim that the HTTP and HTTPS protocols are more of a transport protocol than an application. If someone were to ask what kind of content is expected when looking at an HTTP session, I would have to answer that I have absolutely no idea. If the same question was asked about FTP or SIP, I could explain the protocol and what kind of communication is expected to be seen. The problem is that HTTP is ubiquitous and designed to be a flexible framework to deliver all sorts of content ranging from text to images to video to personal information. Looking at an incomplete list of applications that use HTTP, it is hard to believe that we ever thought that using packet headers and 5-tuples was enough to manage content from a network perspective.

The problem with classifying traffic in today's networks is that many applications use destination TCP port 80 which is reserved for HTTP. In addition, even when HTTP traffic is seen, it could be application/soap+xml manipulating database records, video/mpeg downloaded from a website, or application/x-[application] for anything that a developer wants to create. There is potential for an unlimited number of applications and protocols that can use HTTP as a transport mechanism.

The use of a common application for transport means that the ability to inspect the content payload of the traffic on the network is essential when there is a need to provide differentiated services based on protocol and application. The carrier service providers (CSP) and other entities, such as educational institutions and health care facilities, which have a desire to classify and manage the traffic on their network must deploy a technology and solution that has the ability to inspect and classify the network traffic accurately.

F5 Networks, Inc. | 401 Elliot Avenue West, Seattle, WA 98119 | 888-882-4447 | f5.com

F5 Networks, Inc.
Corporate Headquarters
info@f5.com

F5 Networks
Asia-Pacific
apacinfo@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.com