Get Out Your Crayons. We Need to Color Us Some Bits

Lori MacVittie, 2010-22-07

Those eight bits in the IP header aren't doing much of anything these days, perhaps it's time to put them to work

Back in the early days of bandwidth management, when quality of service and prioritization of traffic were on everyone's minds because we were stuck with low throughput connectivity, there was a brief discussion about the use of IP's TOS (Type of Service) bits as a means to meet specific application performance needs. I say brief because, well, it never really got anywhere. See, even though the creators of the IP specification had looked into the future and provided a technical solution to prioritization of traffic they couldn't have looked into the future and seen the organizational roadblocks to leveraging such a simple but effective method of managing traffic.

The biggest problem was that you could ensure that TOS bits were honored within the organization, but once those packets passed through the organization's boundary of control, i.e. onto the Internet, there was no guarantee or requirement that any other organization honor those bits. And because packets flow through many, many different routers and switches along its very long yet ironically very short travels from datacenter to client, if just one fails to honor the bits then the packets go gray and prioritization is lost.

Hence it was that bandwidth management moved up the stack, with queuing and rate shaping at the transport and application layers of the stack becoming the norm and the "coloring" of TOS bits fell into disuse, like childhood crayons set aside in favor of cool gel pens and clicky-mechanical pencils.

THE NEED STILL EXISTS

Interestingly enough, the need for prioritization and bandwidth management still exists and, in many cases, it could become of paramount importance to the successful implementation of a mature cloud.

It is easy to forget that when you look under the covers of a cloud computing environment that there still exist physical network connections that comprise that environment. Despite the magnitude of virtualization in use at all layers that abstracts the entire infrastructure from its physical implementations, that network is still there. It's massive, it's huge, and it's a spider's web of connectivity.

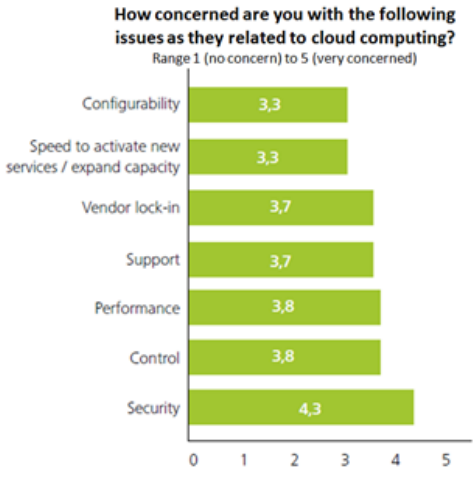

Information Week Analytics Cloud Computing Survey, 2009

I am not the only one thinking about this, I'm fairly certain. An InformationWeek Analytics survey last year included a concern rarely seen in cloud computing surveys: "speed to activate new services/expand capacity." And folks are apparently somewhat concerned about this. Less so than other "problem" areas of cloud but enough concerned that it made the chart.

The problem is that because everyone is still actually sharing a physical network connection all that traffic gets mixed up on the wire. At the physical layer, every packet is the same regardless of what payload it's carrying. But the reality is that some packets are more "important" in the sense that some will be extremely time sensitive. A request to provision a service that's under increasing load is more important than many application requests but both look the same to the routers and switches that get those packets from point A to point B within a cloud computing environment.

As cloud computing providers and enterprises get better at automating the provisioning and elastic scalability processes, the packets that make up requests for such operational tasks will become increasingly time sensitive and important to ensuring operational efficiency and success.

The problem, as once reared its ugly head before, is that all that traffic – customer and operational – is running over HTTP. REST, SOAP, whatever. The APIs that make it possible to provide "compute as a service" are almost unilaterally implemented using a combination of HTTP and REST or SOAP-based architectural principles. And so are the applications running in the cloud. Everything is running over HTTP and thus, as with bandwidth management challenges in the past, it becomes increasingly difficult to distinguish traffic without inspecting its payload.

And payload inspection always adds latency. It may be only microseconds but it's still latency. And if every router has to do it, well, microseconds eventually add up to seconds. And when we're talking about ensuring the order of operations in a provisioning process, those seconds can make a big difference.

BACK to BEGINNING

Sometimes the solution really is to go back to the beginning, to our roots, to the network.

It could be that the solution lies in out-of-band management networks. If not physically at least logically separated network channels that are by default prioritized and therefore will never get stuck in a traffic jam trying to get to location "C" because the application at location "A" is physically on the same network and heavily oversubscribed at the moment an important management request is trying to get to "C". This would, however, add another layer of complexity to the management of not just the physical but the logical network. Complexity that is almost always translated into higher costs, which of course gets passed on to the customer.

It could be that DCSP (DiffServ or Differentiated Services Code Point (DSCP)) – what eventually became of TOS because it was basically unused - could become the solution precisely because a single provider "owns" the entire network. Because a provider has control over all the components comprising the underlying network infrastructure it could enforce the honoring of DSCP bits across its network and thus prioritize traffic at the IP layer.

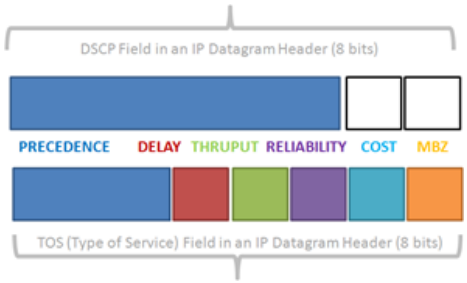

> DiffServ is concerned with classifying packets as they enter the local network. This classification then applies to **Flow** of traffic where a Flow is defined by 5 elements; Source IP address, Destination IP, Source port, Destination port and the transport protocol. A flow that has been classified or marked can then be acted upon by other QoS mechanisms. Multiple flows can therefore be dealt with in a multitude of ways depending on the requirements of each flow. Packets are first Classified according to their current DSCP. Then they are separated into queues where one queue may be routed via a marking mechanism and another queue may be examined more closely.
>
> -- Quality of Service overview

If a provider segments by port "cloud management" traffic from "normal" traffic, this solution would be fairly easy to implement. If not, well, then we're going to need something else. It's the concept behind DiffServ and TOS that's important to leverage in such a solution – the recognition that some traffic must be prioritized to ensure delivery in a timely fashion.

EIGHT BITS

What seems clear, to me at least, is that eventually we're going to run into scenarios in which we need something akin to Operations Performance Management (OPM) to ensure that management and control messages are delivered in a timely fashion.

This is especially true as cloud computing matures and we start to see the dynamism inherent in infrastructure 2.0 components put into broader use. Real-time enforcement of security and delivery policies must be ***real-time.*** They can't be near-time, they must happen ***now***. When we start relying on existing open standards like HTTP and messaging hubs and event-driven networking architectures we have to remember we get the good ***and*** the bad from those existing standards and implementations. We get the ease of use and integration, the flexibility, and the ubiquity of support, but we also get the problems that have plagued quality of service implementations forever: when everything is delivered via HTTP then everything looks like HTTP. Differentiation is nice, but we need to have a way to do that that doesn't impede performance in a cloud computing environment.

In this case, it seems wise to look down the stack and return to a perhaps less sophisticated but absolutely more elegant and simple means of distinguishing not only traffic but precedence and terms of service. DSCP or TOS or whatever we might decide to slap into those 8 bits in the IP header (perhaps there's room for a new specification and use of those bits?) would be infinitely more scalable and easily supported in a cloud computing environment for distinguishing between the small subset of internal "traffic types" that need to be managed.

Related Posts

- The Rise of the Out-of-Band Management Network

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