Deployment Experience with Differentiated Services

Bruce Davie
Cisco Systems, Inc.
1414 Massachusetts Ave.
Boxborough, MA 01719, USA
bsd@cisco.com

ABSTRACT

While ubiquitous QoS mechanisms are not yet deployed widely across the public Internet, the Differentiated Services (diffserv) architecture has in fact proven itself to be a good match for the technical needs of many service providers. In this paper we consider the state of deployment of QoS mechanisms in large service provider IP networks (many of which happen to be offering VPN or VoIP services rather than public Internet service.) We discuss the factors that have helped and hindered the deployment of QoS mechanisms in general and diffserv in particular. We conclude that many if not most of the barriers to QoS deployment are business issues rather than technical shortcomings of the existing QoS architectures.

Categories and Subject Descriptors

C.2.5 [**Computer-Communication Networks**]: Local and Wide Area Networks

GENERAL TERMS

Performance, reliability.

KEYWORDS

Differentiated Services, Quality of Service

ACM SIGCOMM 2003 Workshops, August 25&27, 2003, Karlsruhe, Germany.

1. INTRODUCTION

When the main standards for Differentiated Services (diffserv) were successfully completed[2][6][8][10], it was perhaps to be expected that deployment of diffserv would proceed in the Internet. However, as any user of the Internet knows, deployment of QoS mechanisms in the public Internet remains sparse at best. What is less well known is that diffserv QoS mechanisms have been quite widely deployed by service providers operating large public networks. In this paper we consider the deployment of diffserv mechanisms and attempt to provide some answers to the following questions:

- What factors have led some service providers to deploy diffserv mechanisms?
- What factors have inhibited the deployment of diffserv?
- What factors might encourage the deployment of diffserv, especially in the public Internet?
- Are the diffserv mechanisms well suited to the needs of service providers?

Our main conclusion, based on current deployment of QoS mechanisms, is that the diffserv standards and the implementation of those standards represent a good start to addressing real-world QoS needs.

1.1 Definitions of QoS

When discussing deployment of QoS, it is important to have a clear working definition of what QoS means in this context. There are at least two valid definitions of QoS that might be used:

- QoS may be defined in terms of **application performance**.
- QoS may be defined in terms of **mechanisms** to differentiate the QoS delivered to different traffic classes, such as non-FIFO queueing (e.g., WFQ) or differential drop strategies such as Weighted Random Early Detection (WRED).

In this paper we are primarily interested in the second, mechanism-based definition. This is somewhat paradoxical, because

surely application performance is what really matters. If all networks delivered the low loss, low jitter, and low delay that the most stringent applications require, and used no QoS mechanisms to do so, it is hard to think that would not be a satisfactory result, at least for end users. However, our main concern in this paper is with the deployment of mechanisms, as we seek to address the question of whether the last decade of research, development and standardization of mechanisms has been worth the effort.

2. SERVICE PROVIDER QOS DEPLOYMENT

2.1 Available Diffserv mechanisms

It is worth noting that diffserv capabilities are so widely available on middle- to high-end routing platforms that almost any service provider who chose to implement diffserv in his network could certainly do so, at least as far as router capabilities are concerned. The major router vendors have implemented EF, AF, and class selector (CS) PHBs, as well as a range of "edge" functions such as sophisticated classification functions, token bucket policers, and re-marking functions. Thus the issue is not one of deploying appropriate hardware but rather of investing the effort to configure, provision and manage diffserv capabilities in the providers' networks.

2.2 ISP Edge

While deployment of QoS mechanisms in the *backbones* of ISPs remains rare, diffserv mechanisms are increasingly being enabled on the "edge", i.e. on the access links that connect customers to the ISP. This is typically not a capability that is available to consumers, but for large corporate customers it provides a tool to manage the usage of access links.

There is a strong financial incentive for a customer to buy access links that are appropriately sized rather than over-dimensioned. Assuming that a customer has some traffic that is more important or more sensitive to loss or latency than other traffic, providing the customer with diffserv mechanisms on the access link could enable him to obtain the desired QoS for those more sensitive classes without needing to buy an over-dimensioned access link that would deliver low loss and delay to *all* traffic.

A clear problem arises with this model. Assume that the access link connects a customer-edge (CE) router and a provider edge (PE) router, as shown in Figure 1. Clearly the customer is more easily able to control the diffserv marking, and thus the behavior, of traffic *leaving* his site (i.e. on the outbound direction of his access link) than that of traffic entering his site. He can, for example, mark his own latency-sensitive traffic as EF, but he can't necessarily expect that traffic that he regards as latency sensitive will be marked as EF before it reaches the provider's edge router on its way into his site. One way in which the service provider can help out in this case is to use the source address of the packets as a criterion in determining how the packets should be treated when being sent down the access link to the customer. For example, if the customer connects to the same provider at multiple locations, he could ask the provider to implement a policy in the PE that respects the DSCP value for packets that are sent to this customer if and only if they originated at addresses that are within other sites of the same customer.

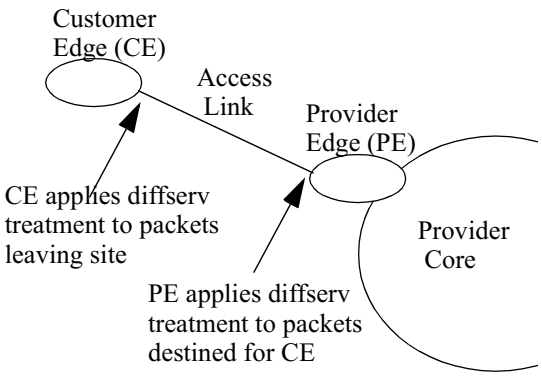

FIGURE 1. Applying diffserv to an access link

2.3 "Layer 3" VPN providers

The discussion of the previous section helps to illustrate why much of the deployment of diffserv today is found in the networks of providers who offer "Layer 3" VPN services. These services offer a virtually private IP network interconnecting the sites of each customer, and are often based on the MPLS/BGP VPN model [12]. Some examples of MPLS/BGP VPN services that offer QoS based on diffserv are discussed in [3][4][7][14]

Because these services connect many sites of a single customer to the network of a single VPN service provider, the problems of offering QoS at the edge of an ISP are largely overcome. The customer can mark his latency-sensitive traffic as EF, his mission critical data with some other DSCP, and so on, and the provider can then honor those markings at the egress edge of the PE routers that connect to other sites of the same customer.

Note that the key difference between the VPN scenario and the ISP scenario is that a single VPN customer is likely to connect to a single VPN service provider at a relatively large number of points, and that the majority of traffic crossing any customer-provider link is likely to be from other customer sites of the same VPN. This makes it possible for the customer to design and enforce his own QoS policies (e.g. packets are only marked EF if they originated at an IP phone) while all the service provider has to do is to honor the EF marking on packets when forwarding them from the PE to the CE (by providing EF behavior to those packets in the customer-facing interfaces of the PE).

Once the customer marking of traffic and service provider treatment of traffic is in place, it is a relatively small step for the VPN provider to start providing diffserv treatment of the customers' packets in the provider's core. The major effort here is in configuring the appropriate PHBs in the core, and typically in configuring policers at the PE ingress interfaces. The latter step is needed to ensure that customers cannot send arbitrary amounts of non-best-effort traffic. Customers in this environment are generally allowed to send up to some token-bucket-limited amount of traffic in each diffserv class.

As a side note, since these layer 3 VPN services typically are based on RFC 2547, the packets in the core are MPLS labelled, and use the procedures defined in RFC 3270[9] to deliver diffserv capabilities in the provider core. From the customer's point of view, there is no difference between an MPLS and a non-MPLS network in terms of its diffserv capabilities - MPLS is used only to provide the constrained connectivity required by the VPN service.

It is worth noting that offering this sort of service also requires the customer and provider to negotiate some sort of "service level specification" (SLS) that defines how much traffic the customer may send in each class and what the provider commits to deliver (in terms of delay, loss etc.) for each class. Such negotiations are commonplace in layer 2 VPN service offerings (such as Frame Relay). Since many of the customers for layer 3 VPN services are migrating from layer 2 services, this may explain why it has been straightforward, and in fact necessary, to include QoS issues in negotiation of the SLAs for these services.

The key point that we wish to stress in this section is that diffserv deployment seems to have succeeded in the VPN context because of some important differences between this environment and the public Internet. The aspects of the VPN environment that seem to have helped diffserv deployment are:

- Each customer connects to the provider in many locations, and thus receives significant benefits from diffserv even when it is applied only to traffic that originates in his own sites and only on the access links
- VPN services are overwhelmingly offered by a single provider, so issues of inter-provider agreements do not arise
- VPN customers are frequently demanding with regard to QoS because of both their prior experience (e.g. leased lines, Frame Relay) and because of the mission critical nature of their traffic.

We will revisit these points when we consider the possible future deployment of diffserv in the public internet.

This is not to say that successful deployment of QoS becomes a trivial matter in the VPN environment. For example, there is still the issue of how the corporate customer enforces consistent policies, e.g. to ensure that only those applications that need low latency service mark packets with EF. There is also the issue of appropriate dimensioning of the access links to handle the possibility that traffic may be sent to one site from all other sites at once, and that EF traffic, for example, may single-handedly overload the link, making diffserv mechanisms quite ineffective. However, it is clear that the barriers to deployment are, at present, low enough in the VPN context for significant deployment to have taken place.

2.4 Public VOIP providers

Another area where diffserv deployment has met with some success is in public voice telephony services run over IP networks. Certainly there have been some IP telephony providers who have simply obtained Internet access for their VOIP gateways from large ISPs running lightly loaded best-effort networks, and the typically low latencies across the networks are sufficient to provide good quality voice services. This is somewhat analogous to the early DARTnet experiences that fed into early work on Integrated Services[5] - a lightly loaded network provides essentially the same service as a lightly loaded priority queue, which can meet the needs of latency- or loss-sensitive applications, provided the utilization of the network remains consistently low. As an example of a network that could provide this sort of behavior, the reported metrics from Sprint's IP backbone for the first quarter of 2003 indicate 0.00% packet loss and delay of less than 50ms in the US portion of their network [13].

It is not clear how many providers of public VOIP services are using diffserv rather than just over-provisioned networks. Telecom Italia's announcement of their VOIP offering [15] is an interesting case in terms of both size and visibility - much of the public phone traffic in Italy is simply being moved onto their IP backbone. The press release makes it reasonably clear that they are using diffserv mechanisms to protect the voice traffic from data. (Once again, RFC3270 mechanisms are used to provide diffserv capabilities while also using MPLS for other things, such as traffic engineering.)

Note that the VOIP scenario, like the VPN case, again sidesteps some of the deployment issues of QoS for public Internet applications. The provider typically owns all the VOIP gateways and can thus enforce a consistent policy (gateways get to mark traffic EF, nothing else does). And the traffic can be kept on the network of a single provider, avoiding any issues of inter-provider agreements.

2.5 Enterprise Diffserv Deployment

While the focus of this paper is on service provider QoS, we note that there has been a moderate amount of diffserv deployment in enterprise networks. One primary driver for this is VOIP, with the IP phones and VOIP gateways marking all traffic that they source as EF, and the routers and switches being configured to treat such traffic with the appropriate PHB. This is by no means the only use of diffserv in enterprises; it is also common to use one of the AF or CS PHBs to assure that certain mission critical data applications get access to some minimum amount of bandwidth no matter what other applications are doing.

As in a VPN environment, enterprise deployment is made easier by the fact that a single administrative entity can enforce common policies to make QoS effective, e.g. by ensuring that only VOIP endpoints mark packets as EF.

Increased deployment of diffserv by enterprises may in fact provide some impetus towards service provider deployment, for at least two reasons:

- if either end users or network administrators go to the trouble to set DSCPs to different values to obtain different levels of QoS inside the enterprise network, then clearly that eases the task of providing different QoS treatment in the service provider network as well;

- users and administrators who are able to obtain a range of different QoS levels inside a single campus are more likely to want to extend that differentiation beyond the campus.

QoS is typically considered to be an end-to-end phenomenon: if one link in the end-to-end path fails to provide it, then the application doesn't receive the QoS it needs. Thus enterprise QoS deployment is important when viewing QoS from the application perspective. However, for the remainder of this paper we focus on the service provider portion of the problem, which is where we observe the greatest barriers to deployment of QoS.

3. FACTORS INHIBITING QOS DEPLOYMENT

In this section we list some of the factors that have inhibited the deployment of QoS mechanisms. Foremost among them must be the alleviation of congestion in many large service provider networks during the rapid build-out of backbones that took place in the late 1990s and into 2000. There is simply no reason to deploy any QoS mechanism if packets never experience significant queuing delay or loss. Many large ISPs today continue to run networks at such low utilization that queuing delay is negligible and packet loss due to buffer overflow is almost unheard of. As long as this remains true, there is no point debating the merits of different QoS mechanisms, as far as ISPs are concerned. As noted above, however, the ISP edge is another story; so too are certain peering points. It is also not true that all ISPs have substantial excess capacity on all links, and we can speculate about whether over-provisioning will continue to be acceptable in the future.

There is no doubt that turning on QoS features carries some costs for providers. Personnel need to be trained to configure QoS features correctly, and to troubleshoot incorrect configurations. QoS features need to be tested before deployment into the network. There may also be some performance impact on routers, particularly when performing the more complex "edge" functions such as microflow classification. All these factors represent reasons for providers not to deploy QoS unless there is are offsetting benefits.

If QoS is deployed, there needs to be some negotiation of the SLS, which implies additional work for the provider. The provider also needs to monitor the network at a finer lever of granularity, e.g. to determine if the latency experienced by EF traffic meets the goals specified in the SLS. If the SLS is written in more stringent terms than would have been the case for a simple best-effort service, than additional provisioning effort may need to be invested to ensure that those more stringent goals can be met reliably.

A factor that may be significantly hampering QoS deployment is the differing goals of customers and providers when end-to-end QoS is considered. Large corporate customers would typically like to be able to obtain end-to-end QoS without being locked into the services of a single provider; a provider, on the other hand, would often prefer to encourage his customers to buy all their service from the one provider. Thus providers do not seem to have strong incentives to reach agreements with other providers to enable end-to-end QoS services. In such an environment, a provider who can say "*all* packets will experience minimal delay, jitter and loss" is arguably offering a more attractive service than one who says "I can assure low loss, jitter and delay for *some* of your traffic if you tell me which traffic it is (by marking) and if you buy this more complex, more expensive service".

Another challenge for providers who wish to offer diffserv-based QoS to customers as part of their service is that there are, as yet, no well-defined *services* associated with diffserv. This was in fact an explicit goal of the diffserv working group - to define per-hop behaviors and an overall architecture but to leave service definition to the service providers. The reasoning was that service providers should be free to define services so that they could differentiate themselves from other providers in a competitive marketplace. However, it is interesting to compare this approach with Frame Relay, where the service definition based on committed information rate (CIR) was uniformly adopted by most providers of the service. Presumably providers found other ways to differentiate themselves, such as reliability, price, or even the quality of the "zero-CIR" service (i.e. the best-effort service.) Certainly, the existence of some common, well-understood service definitions for diffserv networks would make end-to-end services possible without complex multi-provider agreements. It would potentially make the issue of negotiating the QoS-related aspects of an SLS more straightforward if a customer was offered the same set of services from multiple providers.

Finally, as noted above, platform support of Diffserv is typically not an issue for service providers. It may however be an issue in the CE devices (see Figure 1), especially for very low-end users.

4. FACTORS THAT MAY ENCOURAGE QOS DEPLOYMENT

Just as excess backbone capacity has reduced interest in deployment of QoS mechanisms, clearly a subsiding of the "bandwidth glut" would probably increase service provider interest in the same mechanisms. The financial environment now is quite different from that which prevailed during the late 1990s. The following piece of anecdotal evidence may be telling: a large provider used a rule of thumb that it was time to purchase a new link from A-B whenever the utilization of the current link from A-B reached 15%. This approach is apparently becoming less acceptable in a more cost-conscious era.

While it is difficult to predict if or when bandwidth in the core will cease to be abundant, clearly the increasing penetration of broadband access is one factor driving up utilization in the core.

Another factor that is likely to inspire providers to deploy diffserv mechanisms in the core is an increasing focus on high quality of service even during failure scenarios. Evidence of this focus is apparent in work on faster routing convergence (e.g. [1]) and other fast reroute techniques (e.g. [11]). It is obvious that even if links are running at 15% or less when the network is completely healthy, then the load may rise far above that in the event of a single failure. A provider wishing to provide a level of service that is appropriate for voice will therefore have a strong incentive to mark and queue voice traffic separately from data rather than run the network at such low

utilization that even under failure conditions all links remain lightly loaded.

Providers may wish to use diffserv mechanisms simply to isolate one class of traffic from another. For example, TCP traffic tends to be highly bursty, and TCP behaves well under congestion, in the sense that it backs off. Many applications that do not use TCP (e.g. telephony) are both less bursty and less responsive to congestion. Thus, it is potentially good for both classes of traffic if they don't share a queue with the other. TCP traffic benefits by not sharing a queue with non-responsive traffic, and voice benefits by not sharing a queue with bursty traffic that is likely to fill the queue quickly and thus induce jitter. This argument doesn't carry much weight as long as links are truly underloaded. However, one point to note is that a link that is loaded at 15% on average may still exhibit high queueing delays for short periods, and that these may be long enough to impact voice quality.

It is reasonable to expect that, as providers gain experience with QoS mechanisms in a VPN environment, there may be a greater willingness to turn on QoS for non-VPN traffic, in those cases where the same provider offers both VPN and public Internet service. At this point, the cost of training, configuration, and monitoring has already been paid, and so the benefits of enabling QoS mechanisms need not be so large to justify the incremental cost.

Finally, we return to the examples of deployment discussed in Section 2. What would it take to make ISP deployment of diffserv catch up with deployment in VPNs and in VOIP networks? One factor would be strong customer demand for end-to-end QoS across providers. Another would be co-operation among providers to use Diffserv in common ways. This might be more likely with providers who lack global presence, and who therefore have an incentive to deliver QoS to customer sites that are not directly connected to their own backbones. It seems likely that well-defined services that can be offered over diffserv networks might also help in getting inter-provider QoS off the ground.

5. CONCLUSIONS

The deployment of diffserv in the public Internet has not proceeded at the pace that might have been hoped for, and the ability to obtain end-to-end QoS assurances across arbitrary sections of the Internet still seems some way off. However, the fact remains that diffserv *is* now deployed in a large number of service provider networks and appears to be meeting a range of needs well. The biggest reason for lack of deployment of diffserv appears to be excess capacity in the ISP backbones, which can hardly be interpreted as a negative statement about diffserv. If excess capacity existed *everywhere* in the Internet, and was sure to do so forever, it would clearly be time to give up on the idea of deploying any QoS mechanisms in the Internet. When we consider the fact that the Internet includes expensive access circuits from customers to providers and peering points that are not always well provisioned, it seems that that time has not arrived.

In one of its primary goals - scalability - diffserv has been extremely effective. The fact that so many providers have turned on diffserv features in large networks, even if only for VPN or VOIP applications, indicates that the scalability hurdle, which is a high one, has been successfully cleared.

The PHB definitions of diffserv also appear to have been very effective. They are all implemented in a wide range of products and they are all in use in large provider networks today.

This paper has not addressed the deployment of diffserv within corporate (or "enterprise") networks. However, the rising popularity of VOIP in large corporations has been a clear driver for diffserv in those environments, where it appears to be working well.

In most cases, the barriers to wider diffserv deployment are non-technical. That of course does not mean that are unimportant, nor does it mean that they are insoluble. There are also areas where further technical work may help overcome the business issues. For example, a concerted effort to define standard end-to-end services that could be deployed across multiple providers using the existing diffserv mechanisms might provide a useful tool to help ISPs deliver QoS to their customers in a way that meets the customers' needs.

Thus we conclude that diffserv is, on balance, a successful technology. It has not turned out to be the QoS panacea that some proponents may have hoped for, and there is more work to be done to see it widely deployed in the public Internet. However, we believe that it is an effective starting point for QoS in large service provider networks that is likely to see increasing deployment in the future.

Acknowledgements

I wish to thank Anna Charny for her review comments, and the anonymous workshop reviewers for their feedback. Thanks also to Clarence Filsfils for providing a great amount of information on the deployment of diffserv-based SLAs by service providers, and to Amrit Hanspal and Azhar Sayeed for their help in finding diffserv deployment information.

REFERENCES

[1] Alaettinoglu, C., V. Jacobson, and H. Yu, *Towards Millisecond IGP Convergence*, in NANOG Meeting, Washington, DC, October 2000.

[2] Blake, S. et al. *An Architecture for Differentiated Service,* RFC 2475, December 1998.

[3] Cisco Systems, Inc. *Bell Canada Customer Profile.* http://www.cisco.com/warp/public/732/Tech/mpls/docs/BellCanada_CP_0820.pdf

[4] Cisco Systems, Inc. *Infonet Customer Profile*. http://www.cisco.com/warp/public/732/Tech/mpls/docs/infonet_0723.pdf

[5] Clark, D and S. Shenker, and L. Zhang, *Supporting real-time applications in an integrated services packet network: architecture and mechanism*, in SIGCOMM Symposium 1992, Baltimore, Maryland, August 1992, pp. 14-26

[6] Davie, B. et al. *An Expedited Forwarding PHB (Per-Hop Behavior)*, RFC3246, March 2002.

[7] Equant. *MPLS - the cornerstone of Equant's IP service portfolio.* http://www.equant.com/content/xml/connectonline_may03_mpls.xml

[8] Heinanen, J. et al. *Assured Forwarding PHB Group*, RFC2597, June 1999.

[9] Le Faucheur, F. et al. *Multi-Protocol Label Switching (MPLS) Support of Differentiated Services*, RFC3270, May 2002.

[10] Nichols, K. et al. *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers,* RFC 2474, December 1998.

[11] Pan, P. et al. *Fast Reroute Extensions to RSVP-TE for LSP Tunnels,* Work in Progress (Internet Draft), Feb. 2003.

[12] Rosen, E and Y. Rekhter. *BGP/MPLS VPNs*, RFC2547, March 1999.

[13] Sprint. *Network Overview: Metrics.* http://www.sprint-biz.com/about/network_metrics.html

[14] Telkamp, *T. Traffic Characteristics and Network Planning.* http://www.nanog.org/mtg-0210/telkamp.html

[15] News release: *Telecom Italia takes Europe's biggest step in Voice over IP (VoIP),* October 2002, http://news-room.cisco.com/dlls/prod_100902.html