

Interworking Issues

This chapter contains four main parts. The first part discusses interoperability issues among the Class Selector PHB (CS-PHB) group, the Expedited Forwarding PHB (EF-PHB) group, the Assured Forwarding PHB (AF-PHB) group, and the Dynamic RT/NRT PHB (DRT-PHB) group. The second part addresses other IP quality models from the viewpoint of a Differentiated Services framework. The QoS evolution of the Internet will largely depend on the smooth interworking of different technologies, particularly Integrated Services (IntServ), RSVP, MPLS, and Differentiated Services (DiffServ).

To further widen the scope, some non-IP technologies are discussed in section 8.3, “Interworking with Non-IP Networks.” ATM networks are widely used in high-capacity backbone networks. The capability of ATM to offer quality differentiation for IP networks is, however, somewhat questionable because of the differences in service models. A significant part of the Internet traffic goes through local area networks (LANs). Contemporary LANs have limited QoS mechanisms, but IEEE 802.1p could mean a pivotal change in that respect. The QoS capabilities for data services are also under strong development. Section 8.3.3, “Wireless Networks,” discusses some general wireless issues with a closer look at General Packet Radio Service (GPRS). Finally, some issues related to multicasting services are discussed.

8.1 Interworking Among Differentiated Services Models

One of the conjectures of the Differentiated Services effort is that when the working group has designed a number of building blocks, the natural evolution process generates a reasonable combination of network services. Perhaps something like that will happen, although it is hard to identify any network service that has merely emerged without an explicit design process. The prevalent ISP business model based on flat-rate pricing comes perhaps closest to that evolutionary process. But even that is an example of a business model rather than

an example of a network service. Network service based on the best-effort model has been used throughout the history of the Internet.

Regardless of the form of evolution, those service providers that can either predict the next step of evolution or define it have a great advantage compared to other service providers. The evolution of Differentiated Services starts with the selection of PHB groups. “The competition will generally be most severe between the forms which are most like each other in all respects,” Charles Darwin noticed already 150 years ago (Darwin 1972). What does this mean in the case of PHB groups, or the forms of Differentiated Services? At length only a few PHB groups can survive; moreover, each of the surviving PHB groups will have to find a distinct place or application. One niche cannot maintain several species.

This section assesses the prospects of four PHB groups—the Class Selector PHB group, the Expedited Forwarding PHB group, the Assured Forwarding PHB group, and Dynamic RT/NRT PHB group—with respect to the following topics:

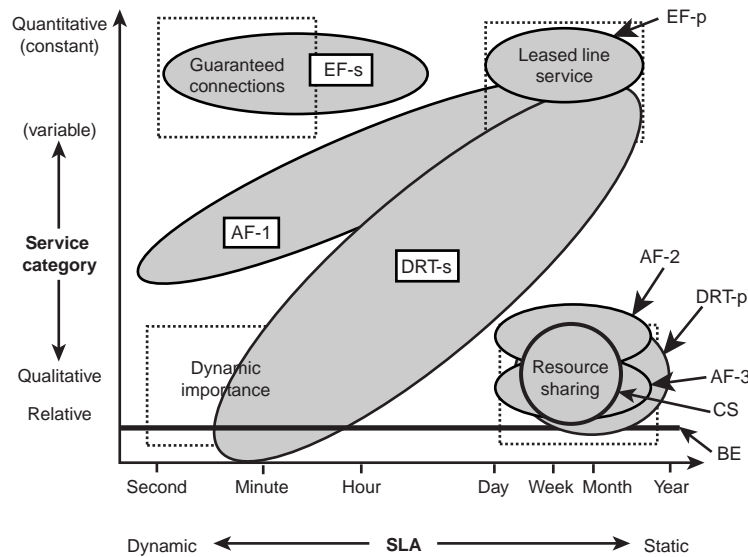
- Class Selector PHB group with other PHB groups
- Expedited Forwarding versus Assured Forwarding
- DRT-PHB group compared with AF- and EF-PHB groups
- Summary of PHB groups

8.1.1 Class Selector PHB Group with Other PHB Groups

The first topic is the general usability of the Class Selector PHB group if the service provider uses one or several other PHB groups. Figure 8.1 shows the applicability areas for different PHBs:

- AF-1 = Main area covered by the highest importance level of Assured Forwarding PHB
- AF-2 = Main area covered by the middle importance level of AF-PHB
- AF-3 = Main area covered by the lowest importance level of AF-PHB
- BE = Best-effort service (default PHB)
- CS = Main area for Class Selector PHB group
- DRT-p = Primary area for DRT-PHB group
- DRT-s = Secondary area for DRT-PHB group
- EF-p = Primary area for Expedited Forwarding PHB group
- EF-s = Secondary area for EF-PHB group

Figure 8.1 Applicability areas for PHB groups.



In the framework of Figure 8.1, there is apparently no empty space to be filled by the Class Selector PHB group. Only if the operator decides to use merely an EF-PHB group is it reasonable to apply a CS-PHB group either for providing better than best-effort service or facilitating traffic management inside the network. The conclusion of Figure 8.2 is very similar: Although a CS-PHB group may enlarge the succinct area of best-effort service, AF and DRT groups do the same job, but better, mainly because of more systematic approach for traffic conditioning in boundary nodes.

Finally, Figure 8.3 shows that AF- and DRT-PHB groups can, at least in principle, cover almost the whole range of importance and urgency levels. It is therefore reasonable to assume that CS is used merely to provide backward compatibility.

262 DIFFERENTIATED SERVICES FOR THE INTERNET

Figure 8.2 Predictability characteristics of the CS-PHB group compared to other PHB groups.

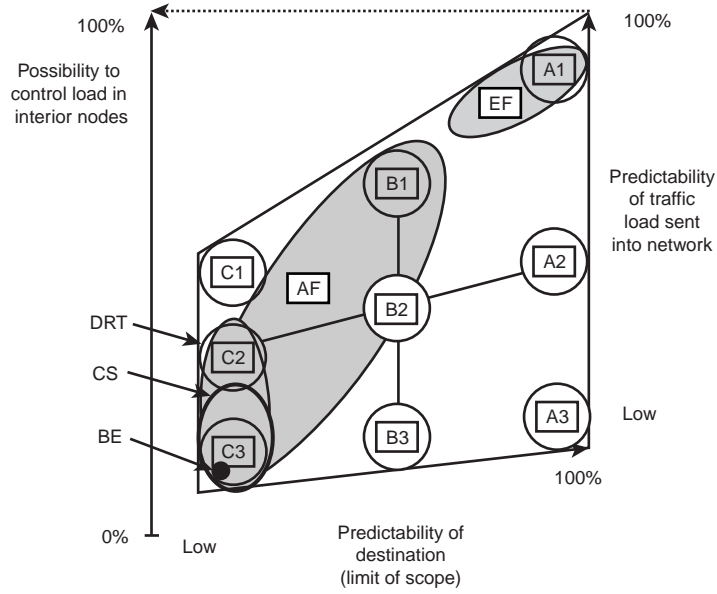
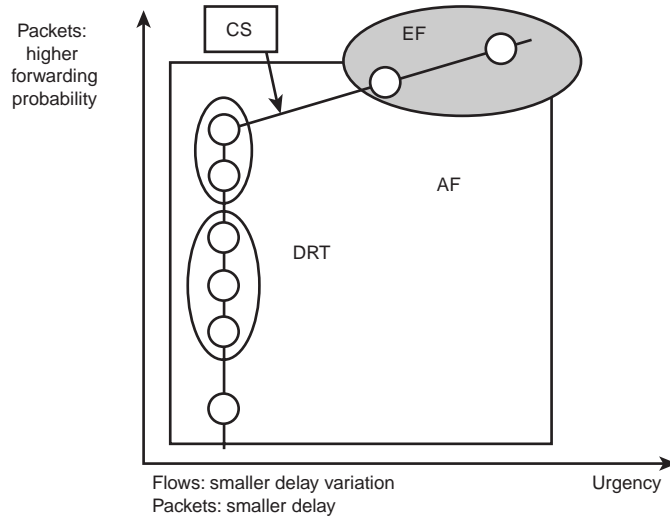


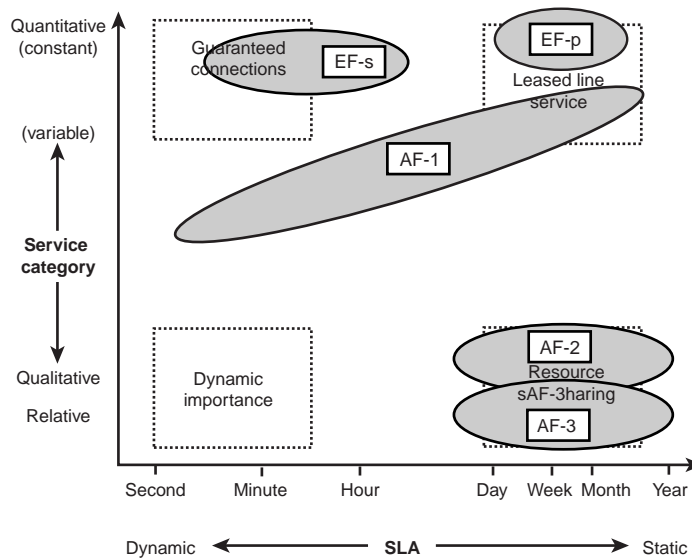
Figure 8.3 The CS-PHB group compared to other PHB groups with regard to importance and urgency scales.



8.1.2 Expedited Forwarding Versus Assured Forwarding

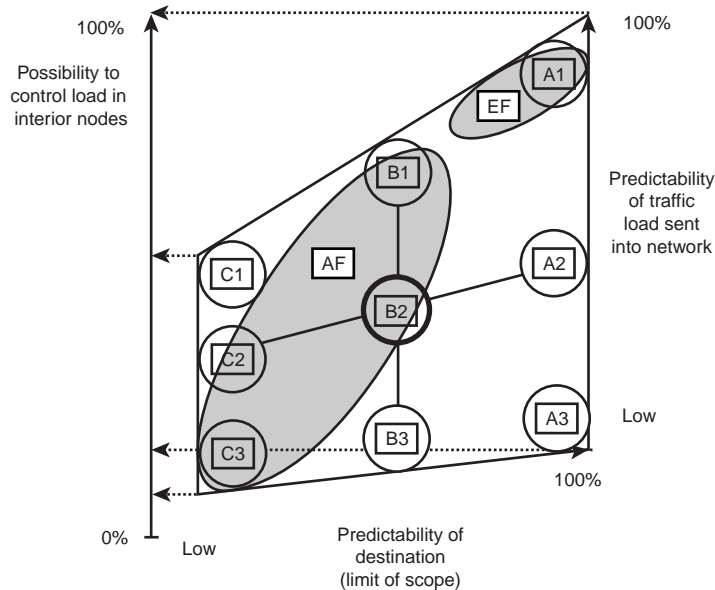
The Expedited Forwarding PHB and the Assured Forwarding PHB group have some clear similarities. The target and implementation of the highest importance level of AF could be similar to those of EF-PHB. Therefore, if both Assured Forwarding and Explicit Forwarding are used in the same network domain, it is important for the operator to identify distinct roles for both PHB groups. In general, Figures 8.4 and 8.5 show that the application regions for EF and AF are mainly separate, which indicates that they can be used in a reasonable manner within one network domain.

Figure 8.4 Service model comparison of the EF-PHB and AF-PHB group.



The operator may, nevertheless, want to minimize the number of PHBs to facilitate network management. The primary idea behind most PHB proposals is to provide specific treatment for packets that have passed certain traffic-conditioning actions at boundary. Now you may ask whether *every* different conditioning action requires its own DS codepoint. The viewpoint of this book is that the logic and structure of the PHB system should be chiefly independent of the boundary functions. That means that a new codepoint should be introduced only if there is a compelling reason, and even then it should be fitted in the existing system of PHBs.

Figure 8.5 Comparison of the EF-PHB and AF-PHB group with regard to importance and urgency scales.



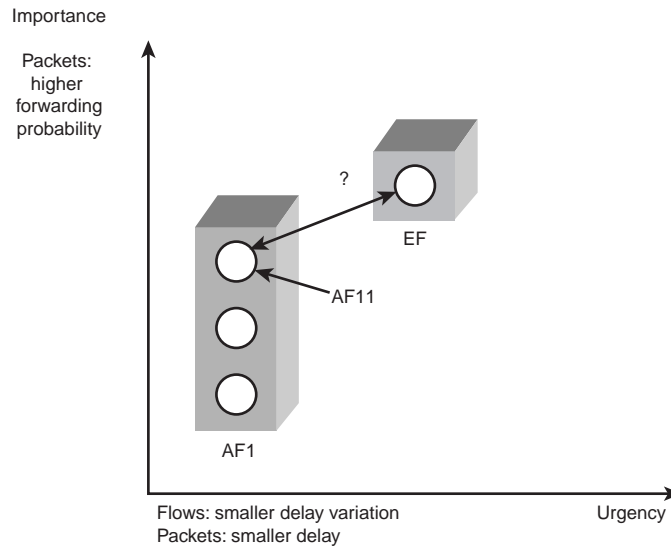
In this case, when you are assessing the interoperability of the EF-PHB and AF-PHB groups, the key question can be derived from Figure 8.5. Is a higher predictability of destination a good enough reason for reserving for EF-PHB its own codepoint? It might be, although the basic idea of both EF and the highest importance level of AF is that the network can somehow make certain that the packet-loss ratio will be minimal. In addition, because the relationship between different AF classes is an open issue, it is hard to map EF into any specific AF class.

It is fair to conclude then that a distinct EF-PHB could be useful. The next issue is to specify the proper relationship between AF and EF. It is easy require that a certain PHB be treated in a certain way (usually very well) regardless of any other traffic streams in the network. But how can anything truly be independent in such an environment as the Internet? An attractive solution to this problem is that the network gives each PHB class certain link and buffer resources. Yet it seems that systematic quality differentiation and strict separation (true independence) are hard to achieve at the same time.

Consider now a system with only one AF class and an EF-PHB, as illustrated in Figure 8.6. From the service provider viewpoint, there has to be a reason for having both AF- and EF-PHBs. In particular, the relationship between EF and highest AF-PHB (AF11) is critical. Should EF be better regarding delay or loss ratio, or perhaps another unidentified

aspect? It is possible to investigate this from the perspective of the fictitious service provider, Fairprofit.

Figure 8.6 The location of an AF class and EF-PHB in the scale of urgency and importance.



Implementing AF- and EF-PHBs in the Same Network

Consider first the loss ratio in a situation where Fairprofit has links with a capacity of 20Mbps for AF- and EF-PHBs together. The managers of Fairprofit have identified four possible models to divide the link capacities inside the network:

- *Model A:* Both AF and EF classes have a 10Mbps capacity without the possibility to utilize any free capacity left unused by the other PHB class.
- *Model B:* This ensures that the available capacity for each of the classes is at least 10Mbps. If one of the classes uses less than 10Mbps, however, the other class can use the remaining capacity.
- *Model C:* AF and EF are sharing the whole capacity of 20Mbps with equal importance for AF and EF packets.
- *Model D:* AF and EF are sharing the whole capacity of 20Mbps in a way that EF packets have higher importance than all AF packets.

The problem encountered by Fairprofit is that despite the incoming EF or AF traffic streams being under tight control on the first link, such as one from node 1 to node 4 in Figure 8.7, it does not necessarily guarantee that there is no conflicting situation in any interior node. Nevertheless, it is expected that overload situations caused by pure EF traffic are very rare.

Figure 8.7 Possible location of overload situation with EF- and AF-PHB groups.

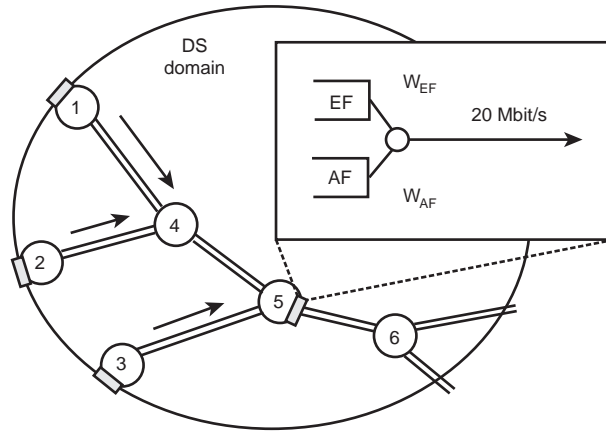


Table 8.2 summarizes the main characteristics of the four models. The table shows the general behavior of the models in five traffic cases, from a case without any overload to a case where both service classes are overloaded.

As long as everything is going well and nothing unexpected happens, there is no discernible difference between AF and EF (no packets are lost). In reality, Case 1 in Table 8.2 should be valid most of the time, say 99.99%, to provide high enough quality. Because there is no difference between different models in Case 1, the only relevant issue is how the system works in exceptional Cases 2, 3, 4, and 5.

Table 8.2 Division of 20Mbps Link Between EF- and AF-PHBs, (For Instance, 12 + 8 Means 12Mbps for EF and 8Mbps for AF)

Case	EF Mbps	AF Mbps	Model A: Strict Separation (at most 10Mbps)	Model B: At Least 10Mbps Available for Both	Model C: Sharing of 20Mbps with Equal Importance	Model D: Sharing of 20Mbps with Higher Importance for EF
1	8	8	8 + 8	8 + 8	8 + 8	8 + 8
2	12	8	10 + 8	12 + 8	12 + 8	12 + 8
3	16	8	10 + 8	12 + 8	13.3 + 6.7	16 + 4
4	8	16	8 + 10	8 + 12	6.7 + 13.3	8 + 12
5	12	12	8 + 10	10 + 10	10 + 10	12 + 8

It seems that there is no clear advantage to applying strict separation (Model A), because it wastes resources in a case where one service class exceeds its capacity while there is free capacity in another class. If there is any justification for using Model A rather than Model B, it might be the assumption that a

looser model encourages users to send excessive traffic in the network. That reasoning seems questionable, particularly with EF-PHB, because end users cannot send any excessive packets into the network. Therefore, any excessive EF traffic inside the network is most likely a management fault rather than a result of end-user misconduct.

Fairprofit may consider Model B as practical because the result in every overload case (3, 4, and 5) is justifiable. Still, Models C and D could be worth of further evaluation. Model C means essentially that the packet-loss ratio is the same for AF and EF, whereas Model D gives a clear advantage for EF in all situations. Which one is the best model depends on the objective of Fairprofit—that is, on the end-to-end service model. That is not the only issue to be assessed, however, because Table 8.2 does not take into account the three importance levels of each AF class. Even though Model B appears to be justified in Case 5, that conclusion is not clear if a significant part of the AF traffic is on the lowest importance level.

Table 8.3 shows four different versions of Case 5. In Case 5a, all AF traffic is marked on the highest importance level; in Case 5d, the AF traffic is marked on the lowest importance level. In Cases 5b and 5c, load is divided between the highest and the lowest importance level. Even though Model B is justifiable in Case 5a, it is much harder to justify it in Cases 5c and 5d. Why should AF13 packets be transmitted rather than EF packets in any case? If the answer is that EF packets always have higher importance than AF13 packets, Fairprofit should use either Model C or Model D. The EF specification mainly promotes the application of Model D.

Table 8.3 Division of 20Mbps Link Among EF, AF11, and AF13 (For Instance, 10 + 4+6 Means 10Mbps for EF, 4Mbps for AF11, and 6Mbps for AF13)

Case	EF Mbps	AF 11 Mbps	AF 13 Mbps	Model B: At least 10Mbps Available for Both	Model C: Sharing of 20Mbps, Importance: EF=AF11>AF13	Model D: Sharing of 20Mbps, Importance: EF>AF11>AF13
5a	12	12	0	10 + 10+0	10 + 10+0	12 + 8+0
5b	12	8	4	10 + 8+2	12 + 8+0	12 + 8+0
5c	12	4	8	10 + 4+6	12 + 4+4	12 + 4+4
5d	12	0	12	10 + 0+10	12 + 0+8	12 + 0+8

Fairprofit should consider also the delay characteristics of the services. If EF requires better delay properties than the best AF class, Model C could be quite hard to implement within the AF framework, whereas both Model B and D are possible. An ultimate version of Model D means that EF can reserve the whole link capacity. But that is against the Assured Forwarding specification requiring that certain buffer and link resources be reserved for each AF class.

If Fairprofit decides that EF can be content with the delay characteristics of AF class 1, EF and AF11 may share the same PHB (Model C in Table 8.3) or share the same class but with a higher importance level for EF (Model D in Table 8.3). Unfortunately, also these models seem to be against the AF and EF specifications because the treatment of EF traffic should be independent of the other traffic streams.

A virtual rescue from this dilemma could be that the behavioral requirements are strictly valid only under reasonable operating conditions and traffic loads. Fairprofit may conclude that because any excess of EF traffic inside the network means unreasonable operating conditions, the specification does not give any

strict rules for those situations. In summary, Fairprofit may consider both Models B and D to be acceptable. Moreover, Fairprofit may claim that traffic control and the network-management systems are able to totally avoid overloads within EF PHB.

8.1.3 DRT-PHB Group Compared with AF- and EF-PHB Groups

The fourth PHB proposal evaluated here is made by the author of this book. This section, however, offers an objective comparison between it and other proposals as far as that is possible. Figures 8.8 and 8.9 illustrate the position of DRT-PHB group compared to EF- and AF-PHB groups. It is important to first concentrate on the relationship between AF and DRT, and return to the EF issue at the end of this section.

Both Figures 8.8 and 8.9 show that the application regions of AF and DRT are similar. For instance, AF drop levels 2 and 3 can be used to implement the resource-sharing model that is the primary area of DRT-PHB group. Therefore, it is reasonable to ask whether there is any good reason to use DRT- and AF-PHBs at the same time.

Figure 8.8 Service model of the DRT-PHB group compared with the EF-PHB and AF-PHB groups.

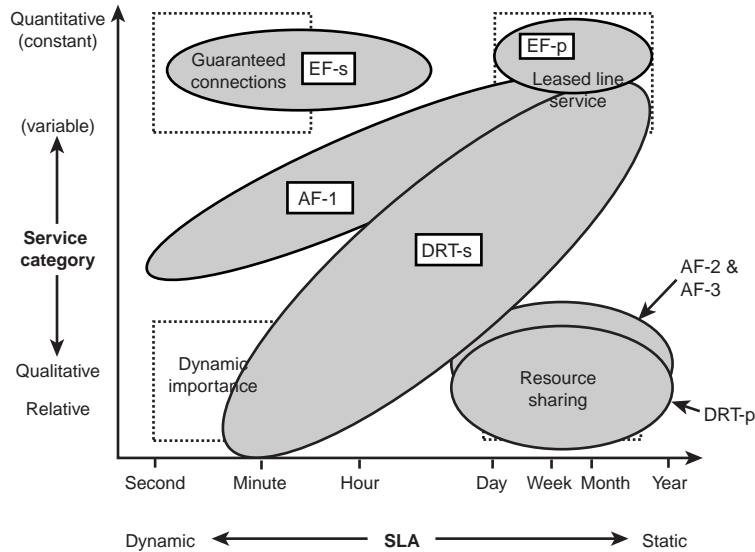
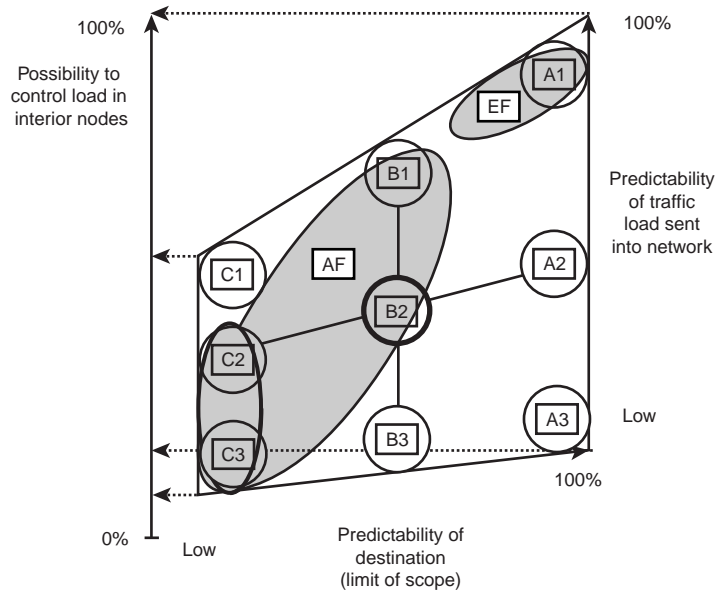


Figure 8.9 The DRT-PHB group compared with EF and AF with regard to predictability of traffic and destination.

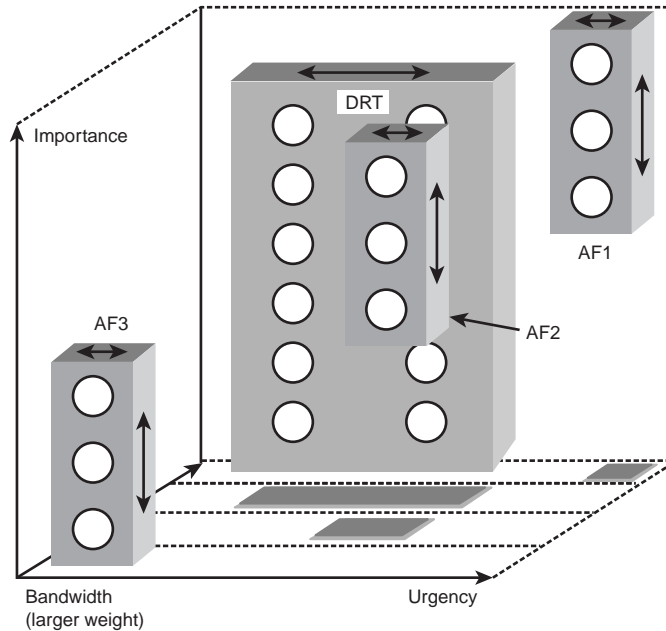


The Assured Forwarding PHB group appears to have the largest area of application. Although this supposition may usually be valid, it is unclear whether the whole area can be covered at the same time by any real implementation. Note that unlike AF the other two PHB groups, EF and DRT, both have their primary implementations that clearly define the application region in Figures 8.8 and 8.9. Hence it is difficult to further evaluate AF without specifying the AF service model or implementation to some degree.

As long as the primary separation between AF-PHB classes is bandwidth (or weight), the actual result related to delay and loss remains unsure. Although a table similar Table 8.3 could be devised for one DRT-PHB group and several AF classes, the result would be too complex. It is better, therefore, to let a figure illustrate the situation. Figure 8.10 presents a system with one DRT-PHB group and three AF-PHB classes. It is supposed that each class has a permanent share of bandwidth (weight) and the relative share is on average larger for AF class 1 than for AF class 2, and so on. Yet, the attribute *on average* means that not much can be said about the instantaneous importance and urgency orders. Although this system can be realized, the design of consistent end-to-end services could be too difficult if both a DRT-PHB class and several AF classes are used in parallel.

270 | DIFFERENTIATED SERVICES FOR THE INTERNET

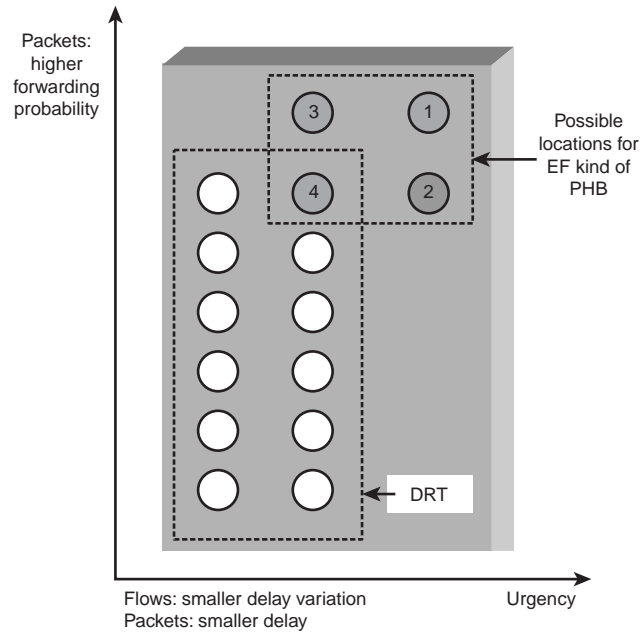
Figure 8.10 Three AF classes (each with three PHBs) and a DRT-PHB group (2*6 PHBs).



As to the combination of EF and DRT, the most logical approach is that an EF kind of PHB be added to the DRT structure. Figure 8.11 shows four possible locations for a virtual EF-PHB. Adjunct *virtual* is used here because the result does not necessarily comply with the EF specification.

Because of the strict requirements of actual EF, it seems that only approach 1 represents an actual EF-PHB. Without its own buffer (approaches 3 and 4) or without distinct importance level (approaches 2 and 4), the characteristics of virtual EF could be inappropriate in some cases. Yet, approach 2 in Figure 8.11 could be suitable for all practical purposes provided that the traffic control can keep the packet-loss ratio of the highest importance level minimal. The main disadvantage of approaches 1 and 2 (with EF having its own buffer) is that they necessitate very strict traffic control for EF traffic to guarantee appropriate delay characteristics for the real-time DRT class.

Figure 8.11 Three AF classes (each with three PHBs) and a DRT-PHB group (2*6 PHBs).



8.1.4 Summary of PHB groups

The main conclusion that can be drawn from this PHB evaluation is that the service provider needs to clearly define the target and then to select a proper system for that purpose. It is unlikely that any rational target will require all PHB groups because the result would be too complex to be managed in a reasonable way. The other conclusions are as follows:

- The main use of CS-PHB can be that the service provider introduces service differentiation by upgrading boundary nodes and applying some CS PHBs in existing interior nodes. In the long run, however, the Class Selector PHB group should be left for backward compatibility, mainly because the system does not provide enough information for proper traffic conditioning in boundary nodes.
- The application of Expedited Forwarding should be limited to cases where both the destination and traffic volume of every EF stream are tightly controlled. Any excess of EF traffic inside the network can be interpreted as an error of traffic management rather than actual misbehavior of any end user. Therefore, even in rare cases of excessive EF traffic, EF packets may have a higher importance than any other PHB used to transmit traffic between end users.

- The main difference between the AF-PHB group and the DRT-PHB group is related to importance order of PHBs. The goal of the DRT-PHB group is to keep the importance order fixed, even between two PHBs belonging to different PHB classes. A proponent of an AF system can claim that the AF structure provides more flexibility in importance ordering; on the other hand, an opponent of AF may claim that it lacks a systematic approach to importance ordering. Because both claims are somewhat justifiable, the final conclusion depends on the service and business models adopted by the service provider.
- Finally, end users encounter the hard task of deciding which one of the numerous service models designed by service providers is the most reasonable for their purpose. Customers should actually be alert because some services are designed more for marketing purposes than for meeting the real needs of customers.

8.2 Interworking with Other Internet Schemes

Differentiated Services is an effort to improve the quality of the Internet service. Yet Differentiated Services cannot just instantly replace the current best-effort network. The evolution toward richer differentiation should be as smooth as possible and in such a way that the users of current services do not experience noticeable degradation of quality. Section 8.2.1, “Best-Effort Service,” addresses this issue.

The first approach to offer advanced quality for the Internet was based on Integrated Services (IntServ) and RSVP. The evolution of this IntServ model has been uneven; the original prospects were promising, but the reality has revealed lots of obstacles in the path toward widespread application. Still the IntServ model has not been vanquished, and there is an inescapable need to define the interoperability between IntServ and DiffServ models. Five primary interoperability models are discussed in section 8.2.2 “Integrated Services and RSVP.”

The third important specification process within IETF related to QoS issues is Multiprotocol Label Switching (MPLS). The original target of MPLS was to facilitate the routing in high-speed networks. The scope of the MPLS working group has, however, also extended to the region of QoS provision. If MPLS will be widely used on the Internet, it is necessary to clearly define its role with regard to the QoS area to avoid needless overlapping of mechanisms and inconsistent QoS models.

8.2.1 Best-Effort Service

It seems that there is one relevant problem regarding best-effort traffic in a network with service differentiation: how to avoid momentary starvation of best-effort service. This may

be considered a fairness issue. To avoid unfairness between “old” best-effort users and “new” DiffServ users, the service level of the current best-effort traffic should not be degraded significantly when new services based on PHBs are introduced.

A service based on EF-PHB should not yield any significant problems because the service provider should keep the load level of EF-PHB relatively low compared to best-effort traffic. The situation could be different with AF- and DRT-PHB groups because in both cases part of the traffic could be quite loosely controlled. In DRT the lowest importance level is intended for lower than best-effort traffic with very few constraints of use, and consequently without any quality guarantees (for instance, during the busiest hours all packets could be lost). The most apparent solution with the DRT-PHB group is that that best-effort traffic is mapped to the DRT system, as discussed in section 7.5.2, “Position of DRT-PHB Group in the Framework.” In chapter 7, the mapping may depend on the load level in the network domain.

As for the Assured Forwarding PHB, a table similar to Table 8.3 can be devised. The main difference between Table 8.3 with EF and AF and Table 8.4 with best-effort and AF is that in the latter case overloads are not unlikely—quite the opposite, both best-effort traffic and the total AF traffic can regularly exceed the reserved capacity. Therefore, the consideration of importance ordering might be more relevant here than in the case of EF-PHBs. Again, for the sake of simplicity this discussion considers only one AF class (AF1) and ignores the middle importance level (AF12).

Table 8.4 Interoperability Models for AF-PHB and Best-Effort Service

Case	AF 11 Mbps	AF 13 Mbps	BE Mbps	Model B: At Least 10Mbps for Both Classes	Model C: Sharing of 20Mbps, Importance: AF11>AF13=BE	Model D: Sharing of 20Mbps, Importance: AF11>AF13>BE
6a	12	0	12	10 + 0 + 10	12 + 0 + 8	12+0 + 8
6b	8	4	12	8 + 2 + 10	8 + 3 + 9	8 + 4 + 8
6c	4	8	12	4 + 6 + 10	4 + 6.4 + 9.6	4 + 8 + 8
6d	0	12	12	0 + 10 + 10	0 + 10 + 10	0 + 12 + 8

Each of the three models shown in Table 8.4 has some drawbacks. Model B appears somewhat questionable in Case 6a because AF11 packets should normally be more important than best-effort packets. On the other side, Model D is not a good solution if the traffic load on the lowest importance level is unlimited. Model C raises implementation problems because equal importance for two PHB belonging to different PHB classes is impossible to

obtain by permanent weights. (Note that if AF class 1 is used for real-time applications, best-effort traffic cannot use the same PHB class.)

One practical approach could be that the service provider applies Model B and keeps the load level of the two highest importance levels below the capacity reserved for the AF class with very high probability. The momentary ordering of the lowest AF packets and best-effort packets is then totally arbitrary. If that result is acceptable from the customer-service viewpoint, however, Model B could be a feasible approach.

8.2.2 Integrated Services and RSVP

Because the authors of “A Framework for Use of RSVP with Diff-Serv Networks” are among the most important ones with regard to the development of DiffServ, that document should be considered as integral specification for Differentiated Services (Baker *et al.* 1998). Although the authors stress that there are different possible scenarios, the only one that they address in detail is a network in which customers use Integrated Services with RSVP for invoking end-to-end service and Differentiated Services is used mainly in the core network (Case 1 in Figure 8.12). The fundamental expectation is that important applications, and users as well, require dynamic high-quality service, and that this demand is the driving force for QoS provision. It is easy to locate this service model in Figure 8.13.

Figure 8.12 Five alternatives for interoperability between Differentiated Services and Integrated Services.

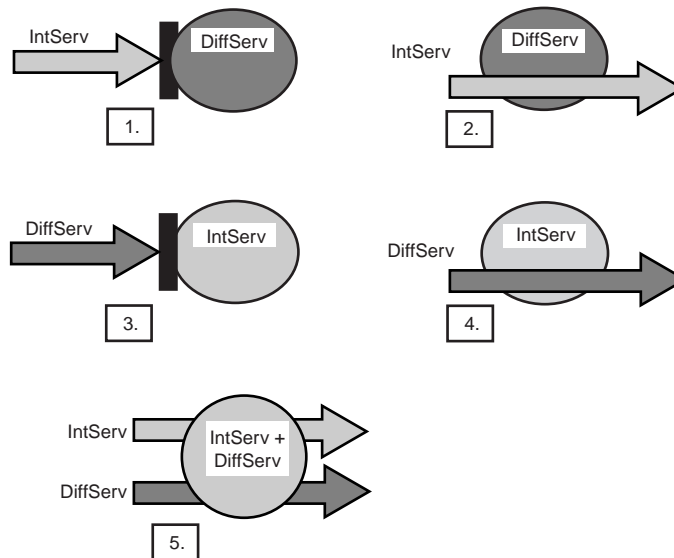
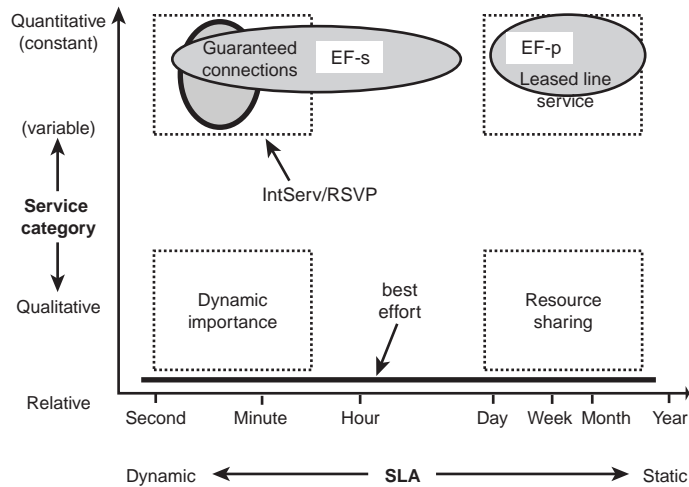


Figure 8.13 The primary service model for Integrated Services.



The opinion presented in “A Framework for Use of RSVP with Diff-Serv Networks” seems to be that the *guaranteed-connection* model is the evident target model, and the main reason why RSVP is not yet used throughout the network is the lack of scalability. The only positive addition of Differentiated Services is that it may provide a looser level of quality—but that is acceptable only because strict quality guarantees are often too difficult to realize. Accepting those opinions, the system described in “A Framework for Use of RSVP with Diff-Serv Networks” makes a lot of sense. To summarize it briefly, a RSVP flow can be implemented by the following seven steps:

1. The sending host generates an RSVP PATH message.
2. In the IntServ network, the PATH message is handled in a normal manner.
3. In the DiffServ network, the PATH message is transmitted transparently through the network.
4. The receiving host generates an RSVP RESV message.
5. In the IntServ network, the RESV message is handled in a normal manner.
6. At the edge between the IntServ and DiffServ networks, the RESV message generates an admission control action in the DiffServ network.
7. If there is enough capacity available, the RESV message is sent toward the sending host.

The critical item is the sixth one, which necessitates an admission control system for the Differentiated Services network. The minimal admission control system consists of a permanent threshold for the PHBs used by Integrated Services. A more complicated system may include the use of bandwidth brokers. An extreme system means that an accurate admission control procedure is made for every RSVP request. (In such a case, a fair question arises: What is left of the simple idea of Differentiated Services?)

As shown in Figure 8.12, there are several other scenarios. Five basic scenarios to combine Differentiated Services and Integrated Services are shown, as follows:

- Integrated Services are used as a customer-service model and Differentiated Services are used in a core network as discussed previously. Integrated Services is mapped to the appropriate PHB within the DiffServ domain. Possible conflicting situations in interior nodes are resolved by means of DiffServ mechanisms.
- The main difference between the first two scenarios is that in the second one the DiffServ network does not provide any statistical multiplexing for IntServ connections. In other words, the DiffServ network is merely used as a transmission medium.
- In the third scenario, Differentiated Services are used as a customer-service model and Integrated Services are used in a core network. Differentiated Services is mapped to the appropriate service within the IntServ domain. Possible conflict situations in interior nodes are resolved by means of Integrated Services mechanisms. This could be problematic, however, because the IntServ model expects that congestion situations are rare core networks.
- The fourth scenario requires tunnels with fixed capacity between each pair of edge nodes, in such a way that a Differentiated Services packet is discarded before it enters the Integrated Services network if necessary. In other words, the IntServ network is merely used as a transmission medium for Differentiated Services. Best-effort traffic can likely be transmitted in a normal manner inside the IntServ domain.
- In the fifth scenario, both Integrated Services and Differentiated Services are implemented throughout the network. This possibility is assessed further in the following paragraphs.

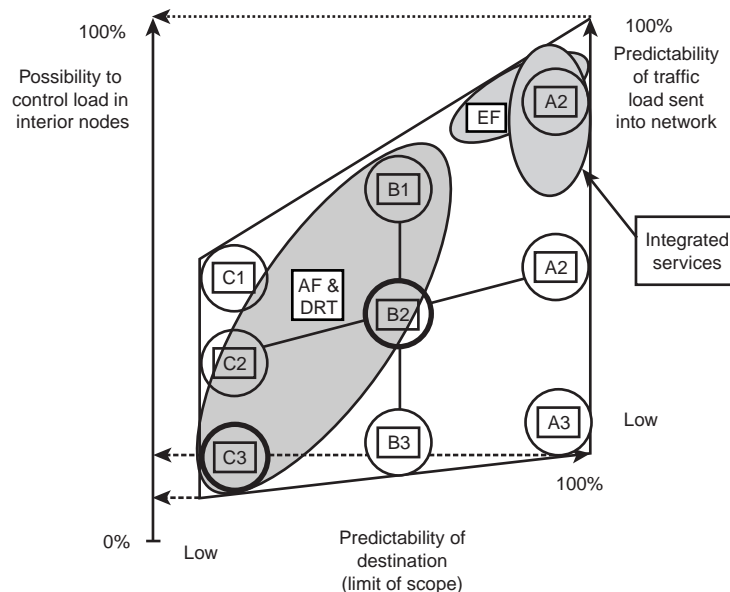
In general, the relationship between EF and Integrated Services is interesting. Figure 8.13 presented the main difference between EF and Integrated Services: EF primarily relies on static reservation, whereas IntServ with RSVP provides dynamic reservations. In contrast, it is somewhat difficult to recognize any significant difference in predictability aspects (see Figure 8.14). EF may, however, allow looser limits for destination according to the principles of Differentiated Services (ingress node does not necessarily know the exact destination of every EF packet). Finally, Integrated Services, particularly controlled-load service, may to a certain extent exploit the variability of the incoming traffic process.

The key issue now is whether these relatively small differences justify clearly different treatment of EF and IntServ packets inside the network. A preliminary answer is that in most cases EF and IntServ packets can share the same PHB inside a DiffServ domain, as long as IntServ flows require real-time service. There are four main reasons for this statement:

- Redundant management and control systems tend to waste some resources, both human and technical.
- Because both services expect an insignificant packet-loss ratio, only very rarely should EF and IntServ packets have to be discerned.
- The importance ordering between EF and IntServ packets is not clear; although the more advanced dynamics of IntServ provides, in theory, better control of the IntServ traffic, the complexity of IntServ may also induce more malfunctions.
- If any kind of statistical multiplexing is utilized, it is better to aggregate as much traffic as possible under one PHB.

Therefore, it is not reasonable to implement specific IntServ mechanisms in a core network using Differentiated Services. Instead, IntServ flows should be mapped to a suitable PHB, such as an Expedited Forwarding PHB.

Figure 8.14 Integrated Services compared with the EF-PHB and AF-PHB groups with regard to predictability.



8.2.3 Multiprotocol Label Switching

The problem statement of the Multiprotocol Label Switching working group (MPLS) identified four relevant problem areas that can be improved by MPLS:

- Scalability of routing
- Flexibility of routing
- Network performance
- Integration of routing and switching

These are very general goals. However, MPLS actually promises to facilitate almost everything stated in “A Framework for Multiprotocol Label Switching”: MPLS makes use of common methods for routing and forwarding over packet and cell media, and potentially allows a common approach to traffic engineering, QoS routing, and other aspects of operation (Callon *et al.* 1997). Although this approach could be successful, history has shown that too ambitious and extensive goals may misguide the development process and conceal the real benefits of the technology.

Because the MPLS working group has recently spent most its effort on traffic management, there is a considerable overlap between DiffServ and MPLS. MPLS has introduced a concept called Forwarding Equivalence Class, for example, which is a group of packets forwarded in the same manner—for instance, over the same path with the same forwarding treatment (Callon, Rosen, and Viswanathan 1997). The purpose of the Forwarding Equivalence Class is basically the same as that of PHB.

Table 8.5 presents a comparison of MPLS, RSVP, and some Differentiated Services approaches to identify the main strengths of MPLS. Routing is clearly an area not covered by RSVP or Differentiated Services. (This discussion does not further address the question about the best way to provide efficient routing in IP networks; there could well be other better approaches than MPLS.)

Network management in general and capacity reservation in particular might be a suitable area for MPLS. If a network operator utilizes MPLS, however, it must be done in close cooperation with other technologies, because it is irrational to make a reservation with several, perhaps inconsistent mechanisms. If any kind of statistical multiplexing is applied inside the network, overlapping reservation systems could be very harmful.

Table 8.5 Targets of MPLS Compared to RSVP and Differentiated Services (X = main target, o = secondary target)

Aspect	Option	MPLS	RSVP	EF	AF	DRT
Facilitating:	Network management	X			o	X
	Routing	X				
Reserving capacity for:	Flows	o	X	o		
	Aggregates	X	o	X	X	
Provided service models:	Guaranteed	o	X	X	o	
	Relative	o			o	X
	Best effort	o			o	o

Building the end-to-end service models on the basis of MPLS does not seem a good idea, because there is no unequivocal MPLS quality model. MPLS might certainly be used with guaranteed and best-effort service, but so far there is no clear insight about general quality or service differentiation. “A Framework for Multiprotocol Label Switching” states that a provision for a class of service (CoS) field in the MPLS header allows multiple service classes within the same label (Callon *et al.* 1997). Later, the authors clarify that the CoS mechanism provides a simple method of segregating flows within a label. The architecture document recognizes that routers may analyze a packet’s header to determine a packet’s “precedence” or “class of service” to apply different discard thresholds or scheduling disciplines to different packets (Callon, Rosen, and Viswanathan 1997). Unfortunately, the document does not elaborate on this issue except to the mention that MPLS allows (but does not require) the precedence or class of service to be fully or partially inferred from the label.

Consequently it could be better to abide by the service rules of Differentiated Services and to try adjusting MPLS to that framework. The recommendation of this book is that if MPLS is used in a DiffServ network, the MPLS label should refer to a PHB class and the CoS field to the importance level of the packet. This statement requires substantiation because it is not necessarily the most prevalent opinion:

MPLS should not deteriorate the service differentiation available without MPLS. Therefore, if MPLS is used for traffic-management purposes inside the network, it should provide basically the same functionality as Differentiated Services.

All packets of a PHB class should be transmitted using exactly the same route and buffer. Hence, different PHBs belonging to one PHB class should use the same MPLS label.

The CoS field is the only available tool for quality differentiation within an MPLS label. If it is not used to inform about the relative importance of the packet, the network node has to look at the IP header for that information. Because of the small size of the CoS field, it can hardly be used for several purposes simultaneously.

One of the main targets of MPLS labels is to facilitate the management of link resources. Because this target is basically the same as that of the AF-PHB classes, it is reasonable to avoid overlapping and inconsistent systems. If both AF and MPLS are used, they should (at a minimum) apply consistent logic.

In summary, the future of MPLS is still unclear, but the main role of MPLS seems to be to facilitate routing and network management on a general level. Because MPLS cannot offer a clear service model, it has to rely on other approaches to provide the service structure. The division of functions between MPLS and Differentiated Services may be based on the principle that MPLS is used to make capacity reservation for aggregate streams and Differentiated Services provides quality differentiation related to delay and loss characteristics.

8.3 Interworking with Non-IP Networks

This concise analysis addresses the most fundamental problems of cases where the implementation of Differentiated Services depends on networking technologies other than IP. Asynchronous Transfer Mode (ATM), IEEE 802.1p, and wireless networks are the technologies addressed in this section. Although this is definitely not a complete list of networking technologies, these three networks provide clearly different viewpoints to QoS issues. ATM has developed for high-speed networks with a large variety of services. IEEE 802.1p is the first real effort to specify quality differentiation for all kind of local area networks. Wireless networks have several specific problems related to the limited bandwidth and the characteristics of a radio channel.

This discussion focuses on the possibility of implementing Differentiated Services in the diverse environment of networking technologies. Unfortunately although many fundamental issues have been identified thus far in this book, perfect solutions to many issues have remained elusive. Remember, however, that the main objective is to identify problems and propose some *possible* solutions.

8.3.1 Asynchronous Transfer Mode

Asynchronous Transfer Mode (ATM) is widely used in the Internet backbone mainly as an intermediary layer between optical transmission and IP packet forwarding. The application of ATM is quite straightforward as long as the IP service is based on best-effort model. The situation is more problematic with Differentiated Services. The original idea of ATM

was crisply based on guaranteed-service models; UBR, ABR, and GFR service categories have been added to solve ATM's inability to offer efficient service for connectionless data applications. Because of this background, it is unlikely that CBR and VBR services can provide an appropriate basis for Differentiated Services apart from PHB intended for (nearly) guaranteed service, such as Expedited Forwarding PHB.

The service model of EF is so close to the CBR model that there is hardly any significant obstacle to using the CBR service to implement EF-PHB. In addition, the UBR service category seems well-suited for implementing best-effort service, as the present Internet has verified. Yet, EF and best effort are not enough to provide real quality differentiation—for real differentiation, something more is needed. This issue is illustrated in the following implementation example of the fictitious service provider Fairprofit.

Implementing AF PHB on the Basis of an ATM Network

Fairprofit wants to use ATM as the core network technology for building advanced Internet service. The targeted service model is based on two AF classes with three importance (drop preference) levels: one class for real-time applications and another one for data applications. The desire is to avoid packet-level processing in the core network, and instead use only high-capacity ATM switches.

First, Fairprofit can consider an approach with a full mesh of virtual paths (VPs) between each node pair. That approach provides only a transport network for IP without any real service differentiation, however, and Fairprofit wants to exploit the network resources and ATM capabilities as efficiently as possible.

Even a larger number of CBR paths, however, do not provide much quality differentiation when the delay and loss characteristics are the same for every path connection. Two quality levels, best effort and high, are certainly more than one level. Nevertheless, this kind of bipolar system does not improve the service offering for a majority of the traffic; best effort is still the same best effort as in the present Internet. Although VBR service allows better statistical multiplexing, the fundamental expectation of VBR service is that cell-loss probability is kept minimal by a proactive traffic-control mechanism. In contrast, VBR service may provide delay differentiation in the form of real-time and non-real-time VBR service categories.

Therefore, to provide equal service for all flows and packets within each AF class, Fairprofit should only use one VP for each AF class between two adjacent network nodes. The reasoning is that packet-loss differentiation within one AF class cannot be built on the basis of several VPs with different theoretical packet-loss ratios. Therefore, Fairprofit has one VP for AF1 and another VP for AF2 on each link. Is this really an implementation of an AF system? I would not say so, although the cell-loss priority (CLP) bit that provides a rudimentary two-level importance differentiation is available. Still, two VBR VPs may provide a moderate solution with the following characteristics:

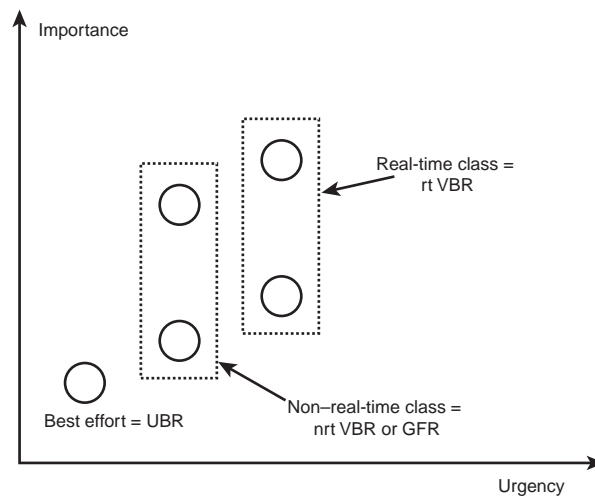
- Two delay classes, the better one suitable for real-time applications.
- Two levels for loss differentiation; that is, the middle and the lowest PHBs use the same CLP marking in the ATM network.
- Cell-loss ratio for the lower CLP marking should (probably) be kept relatively small to provide sufficient quality for the middle AF level.
- Best-effort service should use a separate UBR VP.

282 | DIFFERENTIATED SERVICES FOR THE INTERNET

Figure 8.15 illustrates this system. Fairprofit may suppose that the average load level of real-time service can be maintained lower than that of non-real-time service. Consequently, the relative importance of real-time cells is usually higher than that of non-real-time cells, but this relationship depends essentially on overall traffic management in the network.

Can Fairprofit and its customers be satisfied with this model? Definitely it promises some advantages. But the final result appears something other than a true AF—if anyone really knows what that means. The differentiation provided by ATM seems to be limited to two delay classes and two importance levels. If that is a sufficient service model, Fairprofit should devise and implement a capable network-management and customer-care system to support the guaranteed services.

Figure 8.15 Tentative urgency and importance relationships among rt-VBR, nrt-VBR, and UBR services.



A curious fact is that any service provider has been able to build this kind of system for the past several years using ATM switches. Yet there are not many practical implementations. Although the reason may well be something other than an inferior service model, the lack of concrete use of this model makes it somewhat doubtful. If it is so simple and useful, why there is not yet any real business? The full answer is surely convoluted, but one important aspect is the fundamental service model of ATM that makes real quality differentiation quite difficult to attain.

It is important to finally investigate the issue of what ATM nodes can do with two virtual paths with different packet-loss ratios. It seems that they cannot do anything if they are using the same queue in every network node. The function of the whole system is grounded on the assumption that packet-loss ratio is small, so that there is no need for importance differentiation inside the network. The only clear exemption is the UBR category, which closely resembles best-effort service in IP networks. ABR does not include any

genuine quality differentiation, because packet-loss ratio is supposed to be insignificant and the overall system is primarily intended for non-real-time applications.

The new GFR service category appears to be most similar to Differentiated Services. Actually, when GFR service is available, it is likely to be a more practical choice for data services than nrt-VBR service. From the quality differentiation viewpoint, there is not much difference, however. GFR itself provides two importance levels and only one delay level. Another relevant issue with GFR is the appropriateness of an additional level of best-effort service. The most reasonable approach might be that the lower level of GFR be used to improve the fairness of traditional best-effort service.

GFR is also similar to the service model of Frame Relay, which is comprehensible because the objective of GFR is to offer service similar to Frame Relay. In Frame Relay, the ingress node controls the traffic sent to the network; and if the committed information rate of a flow is exceeded, the node marks packets with lower importance. Frame Relay is also similar to GFR with respect to delay differentiation; there is only one delay class. The congestion-avoidance mechanisms of Frame Relay were discussed in section 6.2.3, “Feedback Information,” in Chapter 6, “Traffic Handling and Network Management.”

8.3.2 IEEE 802.1p

Many end users are attached to a corporate IP network or to the public Internet via a local area network. Traditionally, local area networks, such as the Ethernet, have not supported any quality differentiation. IEEE 802.1p is the first comprehensive effort to extend the scope of local area networks in that respect (802.1p 1998). IEEE 802.1p specification provides a technique to give some packets (formally frames) preferential queuing and access to network resources. IEEE 802.1p offers a consistent method to transport priority information through the local area network regardless of the underlying network layer. To be useful, the network nodes have to be upgraded to understand and utilize the priority bits of packets.

The current specification defines eight levels of user priority, with seven being the highest value. Similar to Differentiated Services, the specification leaves the implementation as open as possible. However, the fundamental assumption is that each level of user priority will have its own queue—in some parts of the document, queue and priority are used almost as synonyms. In contrast, the informative annex of IEEE802.1p does not mention drop thresholds or similar techniques at all. In contrast, the following short evaluation of IEEE 802.1p does not hold to this viewpoint but takes as the starting point the reasonable motivation of each additional priority level. The steps from one priority level upward are as follows:

284 | DIFFERENTIATED SERVICES FOR THE INTERNET

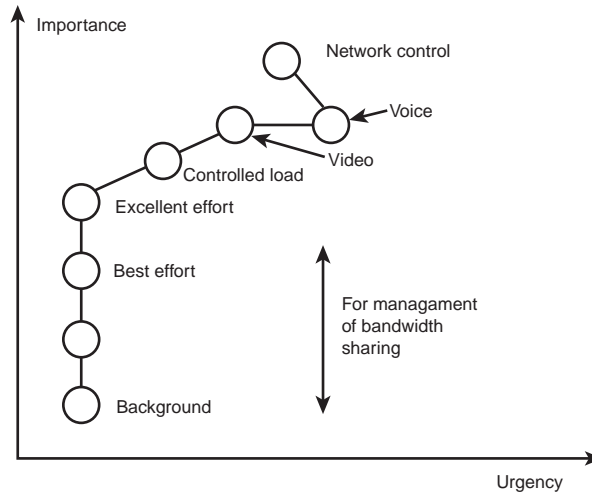
1. Time-critical traffic, such as voice and video, is separated from bursty best-effort traffic.
2. Controlled-load traffic is separated from voice traffic.
3. Best-effort traffic is divided into two parts with different business importance. The lower part is called background traffic.
4. Controlled-load class is divided into two parts in a way that the higher one can be used by video applications.
5. Excellent effort is separated from best-effort traffic.
6. Network-control traffic is separated from voice traffic. It is mentioned, however, that the network-control traffic is probably less delay sensitive than voice but it is of higher importance.
7. The eighth user priority can be used to support bandwidth sharing of data applications.

It is easy to notice that these steps and goals are largely consistent with those of Differentiated Services. There are basically three goals: to separate time-sensitive traffic from bursty data traffic, to provide different importance levels for data traffic, and finally to separate extremely important network-control traffic. If you consider the traffic management of a given network that has altruistic end users, the list of issues is appropriate. The main viewpoint of IEEE 802.1p is the requirements of applications; in Differentiated Services, on the other hand, the main purpose could be to provide building blocks for chargeable services—that is, to facilitate ISP business.

The preceding list of steps makes a good outline for mapping the user-priority system into a scale defined by importance and urgency. Figure 8.16 shows the result of doing just that. Is it possible to map this interpretation of IEEE 802.1p to Differentiated Services? Yes indeed, it might be possible to convert the system to its own PHB group. What is needed is to define reasonable traffic-conditioning actions and clear relationships between the PHBs belonging to the group. A more viable approach is to map the user-priority system to a more general PHB framework.

Another key question is whether there is any significant interoperability problem between IEEE 802.1p and the current PHB proposals. Expedited Forwarding PHB can apparently be mapped into the voice or video user-priority value. A conservative approach is for both voice and video traffic from the local area network to be transmitted in a Differentiated Services network using EF-PHB, and for EF-PHB traffic to be primarily located into voice user priority when transmitted into the local area network.

Figure 8.16 Urgency and importance relationships of IEEE 802.1p.



Assured Forwarding is not as clear a case as EF. A possible interworking scenario with AF and IEEE 802.1p is to use three AF classes, as follows:

- Best PHB class might be used together with voice and network control in a way that network control uses the highest importance level.
- Controlled load and video may share a middle PHB class in a way that video uses the highest importance level, and controlled load uses the two lowest importance levels.
- Lowest PHB class is used to provide a mechanism for bandwidth sharing—that is, for the four lowest levels of user priority. (The two lowest levels of user priority may share the same PHB.)

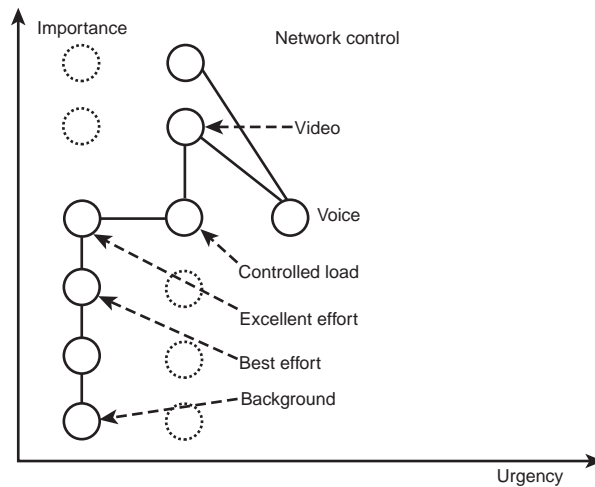
This tentative approach leaves the fourth PHB class for some other use (for instance, for building virtual private networks).

Figure 8.17 presents a possible mapping of IEEE 802.1p priorities to the DRT-PHB group. The mapping of the four lowest priority levels is apparent: They occupy the four lowest importance levels of non-real-time PHB class. The remaining four priorities bring about more difficulties. Controlled load might be on the same importance level as excellent effort, but with better delay properties, or it may share the same PHB as video. The location of video is clear in the sense that it must have both better delay and loss characteristics than excellent effort.

Voice is the hardest matter. It is quite possible in a high-speed core network that voice shares the same delay class as video and controlled-load traffic. Nevertheless, because voice

may need distinct delay characteristics in a low-speed access network, a separate PHB class is reserved for voice in Figure 8.17. Although voice may be very delay sensitive, it does not imply that it is very loss sensitive as well. Actually, the prevalent view is that video is usually more sensitive to loss than voice. A packet-loss ratio of 1% could be acceptable with voice, for example, whereas the same loss ratio is likely to make a video application totally useless. Voice is therefore marked with better delay but lower importance characteristics than video. Network control requires the highest importance level, but it does not necessarily need the best delay class.

Figure 8.17 IEEE 802.1p adjusted into the DRT-PHB structure.



8.3.3 Wireless Networks

The current Internet is chiefly based on wireline links—that is, the network nodes are connected either by electrical or optical cables. Two additional cases are of certain importance and have some special characteristics: satellite links and wireless networks. The most prominent property of satellite links is long transmission delay, particularly if geostationary satellites are used. Although it is reasonable to suppose that real-time flows avoid satellite links whenever possible, they probably do not have any significant effect on the quality model of IP networks.

The situation could be different with wireless networks. The number of wireless terminals with data applications is growing rapidly. Therefore, it is important that the quality model in wireless networks be consistent with that of the Internet. Several aspects make it difficult to reach this target. Certain assumptions that are usually valid in wireline networks—

such as large capacity links and virtually errorless transmission—are not necessarily valid in wireless networks. The main problems of wireless technology relate to the following issues:

- Small bit rates
- Large overheads
- Delay
- Jitter behavior
- Error behavior
- Optimization of packet size

In wireless networks, bit rates vary from 9.6kbps to 20Mbps, which means several orders of magnitude smaller bit rates than in the Internet backbone. Table 8.6 depicts the effect of the bit rates to the transmission delay of one packet. The delay control tends to be very difficult below 64kbps. To attain useful real-time service, the packet sizes should be as small as possible; however, small packet size usually means large overhead. Because bandwidth is a scarce resource in wireless networks, the optimization process related to delay, packet size, and efficiency becomes very complex. The situation is further aggravated by the fact that wireless networks have large overhead because of issues such as encryption information, power-saving signaling, and error-control coding.

Table 8.6 Transmission Delays in Wireless Networks

Packet Size in Bytes	9.6kbps	64kbps	2Mbps	20Mbps
40	33 ms	5.0 ms	0.16 ms	16 ?s
100	83 ms	12.5 ms	0.4 ms	40 ?s
500	417 ms	62.5 ms	2.0 ms	0.2 ms
1500	1250 ms	187.5 ms	6.0 ms	0.6 ms

A further problem is that automatic retransmissions may considerably increase the delay variations or jitter. In addition, because wireless networks are shared mediums, the prediction of delay characteristics is a much more difficult task than in a wireline network where each node has a tight control over every outgoing link.

Consequently, in wireless networks it is practically impossible to give any delay guarantees below 100 milliseconds if the bit rate is, say, less than 100kbps. The only realistic way to considerably improve the situation is to increase the bit rate. Even then, the delay control in wireless networks requires special attention. For instance, it is not reasonable to divide

the relatively small capacity permanently between quality classes. That would deteriorate both the efficiency of statistical multiplexing and the possibility to efficiently manage delays. In addition, because the buffer sizes cannot be as large as in high-speed networks, the queue management may encounter problems when selecting which packets should be discarded (if necessary). Note that the number of incoming packets is much smaller during a given period of time, which makes it more probable that there is no unimportant packet to be discarded.

QoS of General Packet Radio Service

Global System for Mobile communication (GSM) is the most widespread cellular network today. Although the popularity of GSM is based mainly on the versatile telephony service, GSM supports short message service and transmission of circuit-switched data. There is strong demand for developing more efficient service data transmission. General Packet Radio Service (GPRS) is the most important service specification that provides efficient support for typical data applications with bursty and unpredictable traffic. From a Differentiated Services viewpoint, the most essential part of GPRS specification is the QoS profile found in “Digital Cellular Telecommunications System (Phase 2+); General Packet Radio Service (GPRS) Service Description—Stage 2,” which defines the following attributes (GPRS 1998):

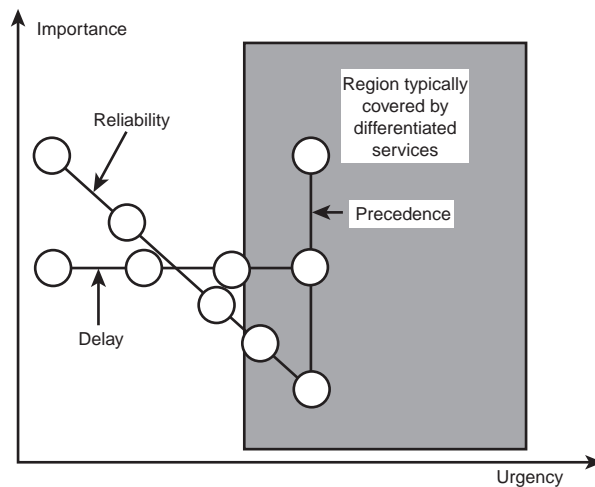
- Peak throughput class
- Mean throughput class
- Precedence class
- Delay class
- Reliability class

Peak throughput class specifies nine values from 8kbps to 2Mbps. Mean throughput class defines 19 values up to 111kbps. The closest points in common with Differentiated Services are the bit-rate values used for traffic-conditioning purposes. The logic of GPRS, however, seems to be that every attribute can be derived from the requirements of the application and that all the attributes are primarily independent of each other. Because there is no traffic-conditioning model in GPRS, the QoS attributes are the same for all packets of a flow. On the contrary, the DiffServ model makes possible a flexible packet marking even within a flow. This discrepancy between the models may yield problems even though the other three attributes—precedence, delay, and reliability—appear to be similar to those used in Differentiated Services.

Unfortunately, a further evaluation of these attributes reveals some additional problems. In principle, three precedence classes of GPRS could be quite similar to those of Assured Forwarding precedence levels if the precedence level of GPRS were not permanent for a flow. Although there are delay classes, there is no class suitable for real-time applications because the delay specifications are upward from 0.5 second. Five reliability classes enable the user or application to define whether the application is more loss sensitive or delay sensitive. The GPRS network uses this information to decide which kind of error-control mechanism is appropriate for the flow.

Figure 8.18 shows one possible interpretation of the GPRS system. In this framework, it is somewhat hard to understand how all three attributes could be orthogonal. In one possible scheme, reliability means that if the network cannot simultaneously fulfill both delay and importance requirements for a flow, the reliability attribute is used to decide how much effort the network should use to transmit the packet at the expense of longer delay. This mechanism relates to the fact that wireless networks cannot offer small packet-loss ratio and small jitter simultaneously because of the fundamental characteristics of radio channels.

Figure 8.18 GPRS quality attributes.



The overall system, with 10,260 ($=3*4*5*9*19$) possible combinations of attributes, is so complex that it is very unlikely that vendors and network operators will be willing to implement the whole range of attributes. Quite the reverse, it is possible that the first GPRS implementations will support only best-effort service without any quality differentiation.

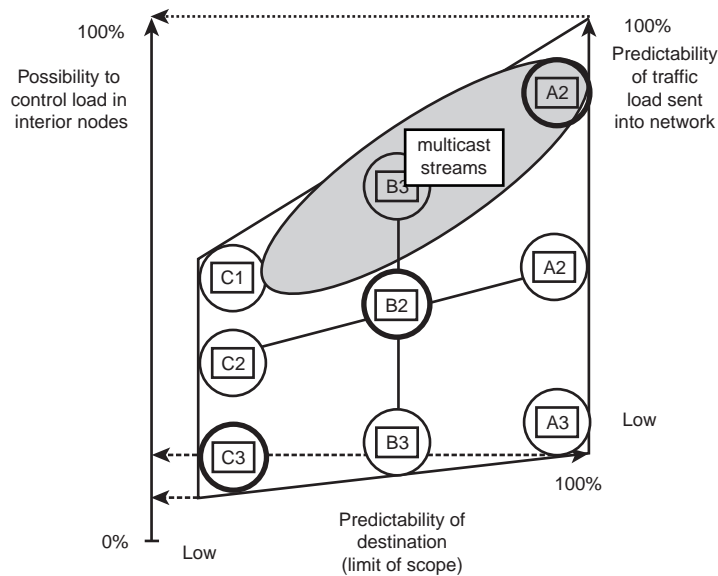
8.4 Multicast Services

Multicast service could be used by various applications, such as streaming public radio stations, software updates, and highly demanding real-time engineering applications. From the DiffServ perspective, two main aspects make multicast special:

- A significant amount of traffic on the future Internet may use multicast services.
- The requirements of traffic management with multicast services could be significantly different from traffic management with ordinary point-to-point flows.

From a technical perspective, multicast streams consume network resources differently than do point-to-point streams. In a Differentiated Services network, the boundary node that classifies and marks packets cannot exactly know how many resources a packet uses later. Therefore, even though the traffic pattern could be constant, the traffic control inside the network is not necessarily able to predict traffic loads on every link. This could be a serious problem if the service model is based on guaranteed quality; with quantitative or relative-service models, however, the situation does not differ significantly from normal non-multicast traffic. Figure 8.19 illustrates this phenomenon.

Figure 8.19 Predictability of load and destination for multicast streams.



As stressed several times, the only relevant situation in a Differentiated Services network is when there are too many packets to be transmitted immediately. Then the system has to be able to make a rational decision about which packets should be discarded or which should be queued. What is the effect of multicast packets to this decision process? This is largely a matter to be resolved by each service provider. It is possible, however, to present some preliminary ideas about the fairness issues related to multicast streams.

Note

It should be stressed that the multicast business and services models are still largely uncertain, which makes it difficult to make any strict evaluation of different technical approaches.

Consider, for example, a service model that primarily takes into account the requirements of each application. Some may think that multicast applications are just used for entertainment purposes and are less personal than applications that use point-to-point connections. (Ignore, for a moment, the fact that these assumptions are apparently incorrect in many cases.) Someone may, for instance, listen to a radio station through the network without paying any attention to the contents of the program. If both *entertaining* multicast packets and *serious* point-to-point packets are present during a congestion situation, we may be inclined to think that multicast packets should be discarded first.

Before concluding too quickly, however, it is important to recognize some significant defects in the preceding reasoning. In public networks, any model that relies on the assessment of applications is prone to raise difficult conflicts. Who can be the final judge for assessing the value of different applications? Even in corporate networks, that would be hard issue; and it is not clear that network-management personnel are always able to make appropriate decisions. Therefore, it might be better to apply a more general approach that is largely independent of the applications using the network services.

The other, perhaps more tractable issue is that although a multicast packet might be judged to not be very important for an individual receiver, there might be thousands of simultaneous receivers. Actually, the number of receivers for one individual packet could vary from one to millions, depending on the location of the packet. In this respect, a multicast packet near the sender could be very important, whereas near the receiver it is unimportant. It is, at least in principal, possible to limit this problem by decreasing the importance level of a multicast packet at some branch points.

One practical question is whether there should be a special PHB or special PHB group for multicast services. Four basic reasons justify this approach:

- Special delay requirements
- Special requirements related to importance marking
- The need to isolate multicast streams from other traffic streams

It is not probable that multicast streams have such unique delay requirements that multicast streams need their own PHB for that purpose. The previously mentioned possibility—to remark multicast packets inside the network—could be a reason to reserve a PHB group for multicast purposes—but the usefulness of this approach is uncertain.

The main justification for a multicast PHB could be to facilitate network management. If the receivers of multicast streams form a distinct customer group, it could be useful to have multicast PHB with fixed resources to protect other users from unexpected quality degradation during high demand of multicast applications. However, the same result might be achieved without a separate PHB group as well.

According to “A Framework for Differentiated Services,” the Differentiated Services architecture deals only with unidirectional flows and therefore each source that wants to send to the multicast group needs a separate SLA (Bernet *et al.* 1998). An additional problem of SLA is that an incoming packet may exit the network domain at multiple points. The contracts between other domains could be different in a way that requires different treatment for basically the same packet going to different destinations. This matter may also promote the use of a special PHB for multicast traffic.

Summary

This chapter addressed interoperability issues on three levels:

- Between PHB groups
- Between DiffServ and other QoS mechanisms for IP networks
- Between DiffServ and QoS mechanisms of other non-IP networks

The main results of the evaluation were as follows:

- Class Selector PHB group does not offer a systematic enough structure to provide an agreeable overall QoS system.
- EF has a clear scope and can usually be integrated with other models, although the mechanisms to solve conflicting situations should be carefully planned.
- There is so much overlapping between Assured Forwarding (AF) and DRT-PHB groups that it is questionable to use them in the same network domain. Note that the status of AF-PHB at IETF is higher because DRT-PHB is proposed only by individual contributors.

- The need for Integrated Services and RSVP depends on the actual demand for highly guaranteed, dynamic connections. If the demand is not considerable, EF or some other PHB should be used in the core networks rather than explicit reservations for every flow.
- Despite the highly promoted QoS mechanisms of ATM, it is a problematic tool when attempting to implement Differentiated Services. All the additional management can hardly be justified by the attained advantages of ATM if Differentiated Services is used as the fundamental service model.
- The role of MPLS is still open, but it seems more suitable for facilitating resource management than for providing quality differentiation.
- The quality models for local area networks and wireless networks are still somewhat unclear, at least when assessed from the viewpoint of Differentiated Services. These issues definitely require more research and development to attain a consistent QoS structure throughout all major packet networks.

