CHAPTER 9

Implementing Differentiated Services

This chapter condenses the entire discussion about Differentiated Services into four implementation examples. The main target is to evaluate as realistically as possible the applicability of the DiffServ models that were introduced in Chapter 7, "Per-Hop Behavior Groups."

Any formal evaluation of an extremely complex issue, such as Differentiated Services, in a real environment is inevitably restricted. If all possible aspects are incorporated into the evaluation, it may turn out that no useful consequences can be made: Certain aspects must be chosen while others have to be ignored. The main aspects evaluated in this chapter are as follows:

- The difference between adaptive and nonadaptive applications
- The difference in relative load levels on different links
- The long-term traffic variations between busy and idle hours
- The efficiency of statistical multiplexing.
- The effects of importance levels

These aspects are investigated in four implementation examples. Each example is analyzed in the same network using essentially the same traffic models. These models are introduced in section 9.1, "Network and Traffic Models," and described in detail in the succeeding four sections. The key elements of the four examples are as follows:

1. In section 9.2, "Improving Fairness by Using an AF-PHB Group," the target is to use one AF-PHB to provide *fair* service between TCP and UDP users within a university environment. An equal service among all end users is preferred, because it is difficult to give every end user a specific service.

2. In section 9.3, “Virtual Private Networks Using an EF-PHB,” a backbone operator provides a *cost-efficient* virtual private network (VPN) service for several large organizations with a number of distributed units. The implementation is based on EF-PHB.
3. In section 9.4, “Service Differentiation with Three AF-PHB Groups,” a service provider wants to offer *versatile* services for residential users. The customer service is based on three levels of quality, called here *grades*. To implement this service structure, the ISP utilizes three AF classes with three importance levels.
4. In section 9.5, “Total Service on the Basis of a DRT-PHB Group,” the three first examples are put into the same *robust* network. That is the real challenge for service providers and network operators. The approach presented here is based on the DRT-PHB group expanded with an EF type of PHB.

9.1 Network and Traffic Models

The global Internet connects tens of millions of computers through a myriad of network domains. There are numerous applications with different characteristics, users with different needs, as well as service providers with different business models. Furthermore, tens of relevant aspects and numerous DiffServ proposals should be assessed. It is impossible to strictly analyze the whole network in any reasonable way with all the DiffServ models. Indeed, the art of mathematical modeling is to make feasible simplifications—simplifications that make the model tractable, but still maintain the essence of the original phenomenon.

The most important simplification made here is that only steady-state behavior is taken into account, whereas all detailed phenomenon are ignored as long as they have no distinct effect on the steady-state behavior. In essence this means that each importance level of each flow is considered as a continuous fluid flow rather than a traffic process of separate packets. Further, every flow is supposed to be constant during a relatively long period and chiefly independent of the perceived quality of the network service. The main expectation is that TCP protocol is supposed to be able to divide the available bandwidth equally between all active TCP flows on a link. Note, however, that in a real heterogeneous network the capacity division provided by TCP is often less fair. Finally, no effort is made to model human behavior when the quality is unsatisfactory.

Clearly, this bunch of simplifications may make the results arguable. But unfortunately there are not too many sound alternatives. An application of an advanced simulation tool might be applied, but hardly with very large number of users. Besides, to obtain essentially more relevant results than in the following examples, human behavior should be modeled as well. That laborious task is left for further study to be done in premium universities and research centers.

Even with a simple steady-state fluid model, innumerable cases and aspects must be studied, including the following:

- Long-term traffic variations
- Significant variations in relative load between links and nodes
- Adaptive and nonadaptive applications
- Several importance levels
- Several levels of aggregation

The following examples address all these aspects.

Figure 9.1 shows the network used in the evaluation. The number of nodes, 14, is a compromise between a realistic, large network and a tractable model. The number of node pairs (91) is large enough to bring about some important aspect related to capacity reservation, but not too large to prevent straightforward analysis. Still for simplicity, the links marked with broken lines in Figure 9.1 are only for protection purposes, not to signify a primary route for any packet.

Figure 9.1 Network structure for implementation examples.

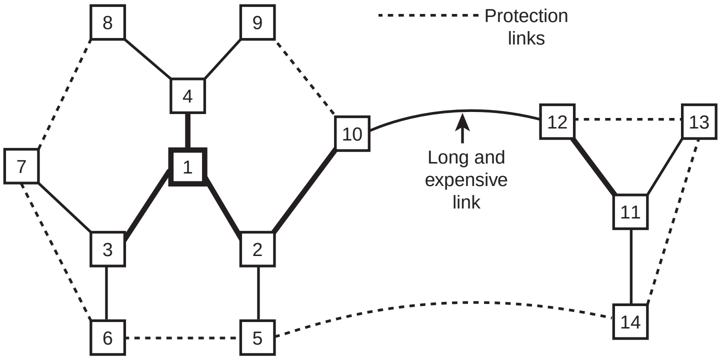

One tricky task is to define a realistic division of traffic between node pairs. Although some kind of randomization could be useful, the approach adopted here (to keep the study comprehensible) is quite systematic and homogeneous. The traffic demand between two nodes is supposed to be proportional to the number of users attached to the nodes. The number of users is constant, but the activity of each user could be different for busy and idle hours. For this discussion, *activity* means the probability that a user is *sending* traffic.

Although this traffic configuration may appear arbitrary, or even unrealistic, a more complex model would make the evaluation very difficult. It is especially important to remember that the whole system should be tractable even in the case of relative complex PHB proposals, such as Assured Forwarding with several PHB classes and importance levels. Apparently, there are various issues not addressed here that need further, more elaborate research. This chapter concentrates on some fundamental issues essential for the realization of viable service differentiation in the Internet.

Assume the following:

Node 1: 10,000 users

Node 2: 20,000 users

All other nodes together: 50,000 users

Probability of activity during busy hour: 0.2

In this case, the number of active users sending traffic from node 1 to node 2 is

0.2*10,000*(20,000 / 80,000) = 500

Note that because users in node 1 are also included in the total number of users (80,000), some of the active users are not sending traffic in the network, but only to other users in the same node. It is also easy to notice that traffic is equal in the other direction from node 2 to node 1.

If an active user sends traffic with a moderate bit rate of 25kbps, the total amount of traffic sent by the user during one hour is 11.25MB. Furthermore, supposing that 20% of the users are active during the busiest hour and that the average load level is double the long-term average load, the amount of traffic sent by an average user is 189MB in one week. Actually, that is equivalent to 6.6 hours of phone calls when a continuous 64kbps coding is applied. Therefore, although 25kbps is definitely not a high value as a peak rate for a flow, it is at least moderate as an average traffic over a longer period of time.

9.2 Improving Fairness Using an AF-PHB Group

The first implementation example evaluates the capability of an AF-PHB group to provide better fairness than a mere best-effort (BE) service. As long as most users are using similar TCP implementations, best-effort service can offer appropriate service for adaptive applications. Still there are some fundamental constraints in the BE service model: It is vulnerable to nonadaptive applications, and an equal share of resources is not always the preferred outcome. This implementation example addresses the first problem, vulnerability; the second problem is discussed in the third example in section 9.4, “Service Differentiation with Three AF-PHB Groups.”

A university environment is used to evaluate this fairness issue concretely. The university is distributed into three sites: 20,000 users in locations near to nodes 3 and 4, and 10,000 users in a location near node 11. A network based on the best-effort model has so far worked well, but now the university has two main concerns:

- The emergence of nonadaptive applications may deprive TCP applications of a significant amount of resources.
- The growth rate of IP traffic has been so high that the university can no longer afford to provide excellent service to the remote site near node 11.

These two issues together require some improvements to the network service model used by the university. Because of the special environment, the preferred service model is still as simple as possible. Specifically, it is not reasonable to assume that each end user is charged based on the network resources he or she has used. The identification of actual end user could be quite a hard task, particularly in the reverse direction. (Note that often the primary direction of traffic is toward the actual user rather than away from the user.) To minimize management costs, each user has to have the same basic rights of use. As a result, the preferred service model is equal service for each individual flow. To meet this target, the university wants to buy better than best-effort service from the service provider Fairprofit (a fictitious ISP introduced in Chapter 1, “The Target of Differentiated Services,” and used in examples throughout this book).

9.2.1 Traffic Model

Concrete information about users and traffic demand is needed before it is possible to make any useful evaluation. The following assumptions regarding the user behavior are made:

- The probability that a user is sending TCP traffic is 20% during busy hour and 5% during idle hour.
- The probability that a user is sending UDP traffic is 0.6% during busy hour and 0.15% during idle hour.
- Each active TCP user is primarily greedy, but the bit rate is properly adjusted based on the lost packets. (For this discussion, *greedy* means that an active user is always sending traffic with as high a bit rate as possible.)
- Each active UDP user is greedy up to 250kbps regardless of the packet-loss ratio.

If the university wanted to offer an available bit rate of 20kbps for each active user during busy hour regardless of the location, the required capacity would be approximately

49Mbps from node 1 to nodes 3 and 4, and 33Mbps from node 1 to node 11. Because the remote link to node 11 is more expensive than other links, however, the university decides to acquire 60Mbps from node 1 to nodes 3 and 4, and only 20Mbps from node 1 to node 11. Note that in principle there is relatively small difference between users in different sites, because the capacity out from node 3 is 3kbps per user compared to 2kbps from node 11. Therefore, the decision to favor users at nodes 3 and 4 seems acceptable.

Because the traffic load is equal in both directions and you are evaluating only a steady-state situation, it is possible to divide each link into two virtual paths (VPs) for analyzing purpose. Figure 9.2 illustrates the bidirectional auxiliary VPs: 50Mbps between nodes 3 and 4, and 10Mbps from node 11 to nodes 3 and 4. It should be stressed that despite these auxiliary VPs, no real VPs with fixed capacity are supposed.

Figure 9.2 Auxiliary VPs between nodes 3, 4, and 11.

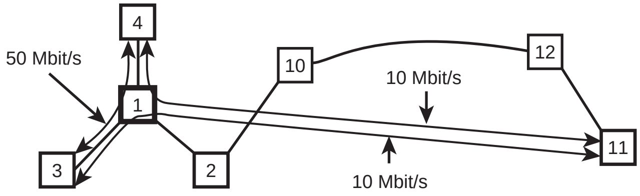

9.2.2 *Implementation*

Fairprofit may consider the following method to implement an equal service for the university, based on one AF class with the following mechanisms:

- A classification unit in boundary node distinguishes every flow.
- A metering unit in boundary node measures every flow by two token bucket devices to define whether the traffic exceeds either of the two predefined thresholds.
- The packet is marked to one of the three importance (drop preference) levels based on the metering result.
- There is only one queue for this service because there is no delay differentiation.
- In every buffer inside the network, packet dropping is based on the importance level of the packet.

In addition, it is possible that the implementation includes the use of random early detection (RED) to improve the performance of the system. This assumption does not have any significant effect on the following performance evaluation.

Because in this case there is only one PHB class, the management of the core networks is probably not much more complicated than the management of a pure best-effort network. Therefore, if appropriate mechanisms are available in all network nodes, the provision of this service could be a relative easy and inexpensive effort for Fairprofit. Correspondingly, the additional price to be paid by the university could acceptable.

9.2.3 Performance Evaluation

This evaluation begins with the best-effort service, to see whether there is any actual problem in the current service. Table 9.1 shows the situation without UPD users. An important observation is that even with moderate differences in link dimensioning and with moderate differences between idle and busy hours, there could be remarkable variations of the available bit rates. Although this is in a way a trivial issue, it is also very easy to forget. If you just consider a separate link with constant traffic demand, these variations related to different links and moments are ignored, and consequently, the results of analysis might be irrelevant.

Table 9.1 Best-Effort Service with TCP Users

Busy / Idle hour	VP from - to	Total Capacity Mbps	Active TCP Users	Capacity per Active TCP User kbps
Busy	4–3	50	1,600	31.3
Busy	11–3	10	800	12.5
Idle	4–3	50	400	125.0
Idle	11–3	10	200	50.0

Next, you add UDP users to evaluate the effect of nonadaptive applications. (For this discussion, the term *UDP* represents all nonadaptive applications.) The number of UDP users is supposed to be only 3% of the number of TCP users, yet they can exploit a significant amount of scarce resources. During busy hour, for example, on average 24 UDP users send traffic from node 11 to node 3. Consequently, UDP applications reserve 6Mbps of the total capacity of 10Mbps. The remaining capacity of 4Mbps is divided between 800 TCP users, which means 5kbps for each TCP flow while each UDP flow attains 250kbps.

Furthermore, remember that this evaluation does not take into account short-term variations. If the traffic load of UDP applications varies significantly, the available capacity for TCP users could occasionally be very small, as Table 9.2 shows. Even the average capacity available for TCP users can vary remarkably, however, while UDP users attain the same bit rate. This is the main issue that should be improved by a more advanced service model.

Table 9.2 Best-Effort Service with TCP and UDP Users

Busy/ Idle Hour	VP	Mbps	Active UDP Users	UDP Load Mbps	Capacity for TCP Users	Active TCP Users	Per Active TCP User
Busy	4-3	50	48	12	38	1,600	23.8
Busy	11-3	10	24	6	4	800	5.0
Idle	4-3	50	12	3	47	400	117.5
Idle	11-3	10	6	1.5	8.5	200	42.5

The first phase of planning an AF-PHB system is to decide the bit-rate thresholds for the two highest importance levels. According to Table 9.1, 10kbps could be assured for all flows and 30kbps could be assured for flows between nodes 3 and 4. Therefore, a packet could be marked according to following rules.

- If there are enough free tokens in the AF11 bucket with a token rate of 10kbps, the packet is marked with the codepoint of AF11.
- Otherwise, if there are enough free tokens in the AF12 bucket with a token rate of 30kbps, the packet is marked with the codepoint of AF12.
- Otherwise, the packet is marked with the codepoint of AF13.

Consequently, each flow is allowed to send AF11 packets with a bit rate of 10kbps and A12 packets with a bit rate of 20kbps.

Tables 9.3 and 9.4 present the analysis of the AF-PHB system. All packets with AF11 can be transmitted successfully. The capacity left for AF12 and AF13 varies from 1.8Mbps to 45.9Mbps. During busy hour, for instance, 1,648 active users are on the auxiliary VP between nodes 3 and 4, which means an average AF11 load of 16.5Mbps and 33.5Mbps for other importance levels.

Table 9.3 Capacity Used by AF11-PHB

Busy/ Idle Hour	VP	Cap. Mbps	Active TCP Users	Active UDP Users	Total AF11 Load Mbps	Capacity Left for AF12 and AF13, Mbps
Busy	4-3	50	1,600	48	16.5	33.5
Busy	11-3	10	800	24	8.2	1.8
Idle	4-3	50	400	12	4.1	45.9
Idle	11-3	10	200	6	2.1	7.9

The remaining capacity is used first by UDP traffic marked with AF12, and then by TCP AF12 traffic as much as there is free capacity. Table 9.4 shows the results. It turns out that all AF12 traffic can be transmitted successfully except between nodes 3 and 11 during busy hour. In that case, you may assume that UDP AF12 packets are transmitted with a bit rate of 20kbps, whereas TCP flows adjust their bit rate down to 1.6kbps to fill up the link. In practice, some UDP packets will be lost and TCP will probably not be able to fill the whole link, but these issues have only a minor effect on the overall conclusion.

Note that UDP traffic on the two highest importance levels cannot fill the whole link in this case because there are not enough active UDP users. As a result, TCP users also get a small amount of resources even though UDP and TCP flows are competing for the same resources.

A similar evaluation can be made for AF13 traffic. The result is that UDP flows can attain the preferred bit rate during idle hours, whereas the available bit rate will be only a fraction of the preferred bit rate during busy hours.

Table 9.4 Capacity Division Between UDP and TCP Flows

Busy/ Idle Hour	VP	AF12 per UDP User kbps	AF12 per TCP User kbps	Capacity for AF13 Mbps	AF13 per UDP User kbps	AF13 per TCP User kbps
Busy	4-3	20	20	0.6	11.7	0
Busy	11-3	20	1.6	0	0	0
Idle	4-3	20	20	37.6	220	87.5
Idle	11-3	20	20	3.8	220	12.5

Table 9.5 summarizes the results of the evaluation. The main conclusions are as follows:

- The results are very promising for busy-hour traffic in the sense that TCP users can attain almost an equal share of resources (99% and 96%).
- Somewhat surprisingly, the AF system has no effect at all during idle hours.

The explanation for the second item is that the bit-rate thresholds were optimized for busy-hour use, and were, therefore, too low to be effective during idle hours.

Table 9.5 Comparison of Best-Effort and AF Services

Busy/ Idle Hour	VP	Equal Share kbps	BE TCP kbps	BE UDP kbps	AF TCP kbps	AF UDP kbps	AF Real/ Equal for TCP, %
Busy	4-3	30.3	23.8	250	30.0	41.7	99
Busy	11-3	12.1	5.0	250	11.6	30	96
Idle	4-3	121.4	117.5	250	117.5	250	97
Idle	11-3	48.3	42.5	250	42.5	250	88

It is important, however, to be somewhat cautious with the results. In general, the selection of bit-rate thresholds is a troublesome task—in practice, there can hardly be as complete knowledge about traffic levels as was supposed in this example. In addition, short-term variations certainly disturb this elegant model. With realistic traffic with strong variations on all timescales, for instance, the bit-rate threshold for AF11 should be lower than 10kbps to guarantee a small packet-loss ratio.

Furthermore, you may ask how those UDP users behave—the ones trying to use 250kbps, but getting only a fraction of the desired bit rate. The answer evidently depends on the characteristics of applications. There are two main alternatives:

- UDP is used by a nonadaptive application with a specific bit-rate requirement.
- UDP is used by an adaptive application merely for getting more resources than with ordinary TCP.

In the first case, the use of AF-PHB system actually yields a high call-blocking probability during busy hours, and may effectively prevent the use of high bit-rate UDP applications during busy hours. This may or may not be an acceptable result. In the latter case with the questionable use of UDP, the effect of an AF-PHB system certainly is appropriate because well-behaved TCP users can get a significant part of the resources deprived by UDP users.

9.2.4 *Possible Improvements*

There are some evident possibilities to improve the overall service model based on an AF class. The bit-rate thresholds could depend on several issues such as time of day, destination, or user. If a user definitely needs connection with 100kbps, for instance, that can be implemented merely by changing the bit-rate threshold for AF11. The main difficulty of this approach is related to the management of rights and bit rates, which probably limits the applicability of this approach to relatively rare cases.

Another possibility is to improve the performance of the AF-PHB system to dynamically adjust thresholds of AF11 and AF12 according to the general load level in the network, or according to the destination of the packets. Both approaches may improve the capacity division between adaptive and nonadaptive applications under variable conditions. Nevertheless, it is not clear whether the attainable gain is large enough to justify the required additional mechanisms and management effort.

Finally, it is possible that the service provider merely has enough capacity in the core network to transmit all packets on every importance level. In that case, PHB classes and importance levels only have a significant effect during exceptional situations, and perhaps in access networks. Sometimes this is the most effective solution for the service provider.

9.3 Virtual Private Networks by Using an EF-PHB

The starting point of this implementation example is the fact that current leased lines are inefficiently used because of poor utilization of statistical multiplexing. The Internet may provide a good possibility to improve the situation, but only if there are proper mechanisms for provision of a service with high quality and reliability. The prevalent best-effort service is apparently insufficient, whereas Expedited Forwarding PHB is designed for that purpose.

The objective of Fairprofit is to provide high-quality virtual private network (VPN) service with low delay and virtually no packet losses for large, demanding customers. Because of this goal, the evaluation is mainly related to network dimensioning made by the service provider rather than traffic analysis. The fundamental assumption is that network dimensioning can keep congestion situations inside the network extremely rare.

9.3.1 Traffic Model

Fairprofit has four big customers. Each of the customer organizations has one large unit and five small units, as shown in Figure 9.3. The following figures defines the traffic demand:

- Each large unit comprises 25,000 employees.
- Each small unit comprises 5,000 employees.
- The activity level of each user is 10%. (Only busy hour is addressed in this example because network dimensioning is based on busy hour traffic.)
- Each active user generates traffic with a bit rate of 25kbps.

The total traffic generated by all 20,000 employees is 500Mbps, but only 350Mbps is actually transmitted within the network, because a part of the traffic is transmitted within each unit. The total traffic load generated by the users appears to be small enough to be handled by a relatively small system without any significant problems.

Figure 9.3 The units of four large organizations.

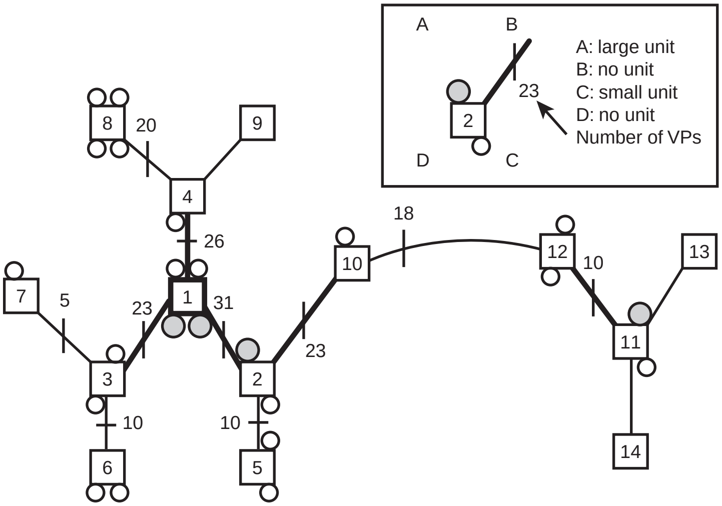

If a virtual path (VP) is established between each pair of units, there are 15 VPs in total for every organization. (Note that *VP* is used here as a general term for any aggregate that is policed as an indivisible entity without any reference to ATM technology.) The number of VPs on each link also depends on the number of hops of each VP. In this example, the number of VPs on a link varies from 5 to 31. Table 9.6 presents the number of VPs and average loads. Note that the last row shows the numbers in one direction, and that the total capacity figure can be obtained by multiplying those figures by two.

Table 9.6 Number of VPs and Load Levels on Used Links

Link	VPs					Load	Mbps			
	A	B	C	D	Total	A	B	C	D	Total
1-2	8	9	9	5	31	30.0	26.3	26.3	11.3	93.8
1-3	5	5	5	8	23	11.3	11.3	11.3	20.0	53.8
1-4	8	5	5	8	26	20.0	11.3	11.3	20.0	62.5
2-5	0	5	5	0	10	0.0	11.3	11.3	0.0	22.5
2-10	5	8	5	5	23	11.3	30.0	11.3	11.3	63.8
3-6	0	0	5	5	10	0.0	0.0	11.3	11.3	22.5
3-7	5	0	0	0	5	11.3	0.0	0.0	0.0	11.3
4-8	5	5	5	5	20	11.3	11.3	11.3	11.3	45.0
10-12	0	8	5	5	18	0.0	30.0	11.3	11.3	52.5
11-12	0	5	5	0	10	0.0	31.3	11.3	0.0	42.5
Total	36	50	49	41	176	95.0	162.5	116.3	96.3	470

On the link from node 1 to node 4, for example, both organizations A and D need eight VPs and organization B and organization C need five VPs. The average traffic transmitted by each of the total 26 VPs is either 6.25Mbps or 1.25Mbps. (In real networks, much more bit-rate variability among VPs should be expected.)

9.3.2 Implementation

The basic way to implement this VPN service is *Expedited Forwarding* (*EF*) PHB. The main property of EF from the viewpoint of this evaluation is that the bit rates are strictly policed. In other words, if there are exceeding packets they are dropped immediately regardless of the load situation inside the network. This characteristic makes it necessary to carefully dimension the capacities reserved for the service.

Now Fairprofit has three primary options to manage and control the VPNs:

- *No sharing*: There is an EF VP between every node pair with fixed capacity for every organization.
- *Partial sharing*: An EF VP is reserved for every organization on every link.
- *Total sharing*: There is only one EF VP for all organizations on every link.

No sharing is equivalent to leased-line service, and therefore quite a clear service model. There is a maximum bit rate from every unit to every other unit of the same organization. In *partial sharing*, each organization can better utilize the bought capacity. If an organization has four VPs of 10Mbps capacity on a link, for example, *partial sharing* means that the total 40Mbps is available for all traffic flows of the organization.

Partial sharing may induce problems inside the network, because the traffic load of an interior link may exceed the reserved capacity even though the load levels in ingress nodes are acceptable. Because EF-PHB does not provide any specific tool to solve this kind of situation, the capacity dimensioning should be quite conservative.

Total sharing provides the most efficient statistical multiplexing. The model, if applied literally, includes an inherent problem. If the traffic control is related only to the aggregate EF stream of all organizations, there is a risk that some organizations may attempt to exploit the situation by reserving less capacity than what they really need. Therefore, a more likely model is one in which the traffic sent by each organization into the network is strictly policed, even though there is only one EF stream on each link. In that case, the main difference between partial sharing and total sharing is in the network dimensioning. In the total-sharing model, the network operator utilizes network resources more efficiently by allowing statistical multiplexing between organizations inside the network.

Figure 9.4 further illustrates the situation. *No sharing* means that each small box (such as the 6.25 in the upper-left corner) is dimensioned separately. An average load level of 6.25Mbps may require a capacity of 30Mbps. *Partial sharing* means that each row with several small boxes is dimensioned separately. Finally, *total sharing* means that the whole area is dimensioned together without taking into account the detailed structure with rows and boxes. In practice, the operator adds up all traffic from all organizations and defines the required capacity based on the total average load and variance.

Figure 9.4 Average bit rates of 26 VPs on link 1-4.

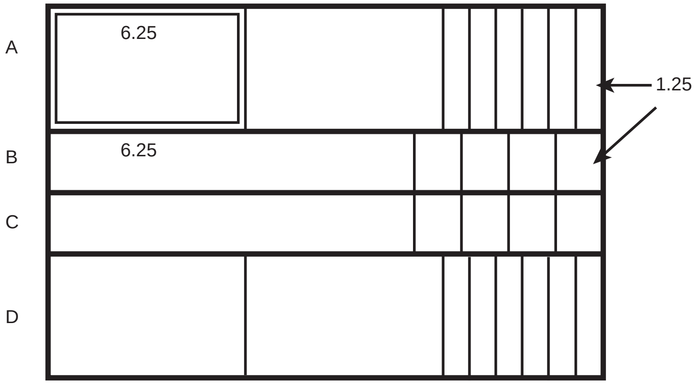

9.3.3 Performance Evaluation

How much does the actual result depend on the sharing principle? To answer this question, it is necessary to define the dimensioning principle applied by Fairprofit. As discussed in section 6.3.3, “Network Dimensioning,” in Chapter 6, “Traffic Handling and Network Management,” a simple but still feasible model is to define the required capacity (C) as a function of average load (A) and variance of the load distribution (V), as shown in Formula 9.1.

Formula 9.1

C = A + ?*V^0.5

If the average load measured on a link is 100Mbps and the corresponding amount of traffic variation (formally, standard deviation) is 20Mbps, for example, the operator needs to have much more capacity than 100Mbps to satisfy customers’ requirements. The parameter ψ defines the level of assurance that the capacity is sufficient to cope with momentary traffic peaks.

Now a good estimation for the average load has been arrived at, but not a practical estimation for the variance. A mathematically elegant approach would be to calculate the variance based on the activity of individual users and the bit rate used by an active user. If there are 2,000 users with a bit rate of 25kbps and activity of 0.2, the average load is A = 10Mbps and the theoretical variance is $(0.2\text{Mbps})^2$. However, that model may significantly underestimate the real variance. Particularly, the greedy and possibly synchronized TCP flows tend to increase traffic variations in real networks. This example adopts a theoretically less elegant, but more realistic model in which the variance is supposed to be as high as V = A*(1Mbps). Table 9.7 shows the required capacity for each link with the three sharing principles, with ? = 10.

Table 9.7 Required Link Capacity with Different Sharing Principles

Link	Offered Load Mbps	No Sharing	Partial Sharing	Total Sharing
1-2	94	592	285	191
1-3	54	380	199	127
1-4	63	436	219	142
2-5	23	162	90	70
2-10	64	418	219	144
3-6	23	162	90	70
3-7	11	81	45	45
4-8	45	324	179	112
10-12	53	337	174	125
11-12	43	237	132	108
Total	470	3129	1631	1132

The differences between sharing principles are apparent. Without any sharing, the load level would be as low as 15%; total sharing raises the load level up to 41.5%.

In practice the differences could be somewhat smaller. The first two alternatives basically leave the dimensioning task for the customers: Customers buy capacity based on the estimation they have about traffic demand in the future. They may use a similar method as was used in this example, but that is not necessarily the case. Some other methods may lead to a smaller difference between no sharing and partial sharing. As an extreme case, the operator decides that the expected EF load should exceed, say 20%, of the reserved capacity. Then there is no difference at all between the sharing principles from a dimensioning point of view.

Nonetheless, from the customer viewpoint there is usually no reason to prefer the no-sharing to the partial-sharing principle, particularly if the capacity dimensioning is conservative. Only if a unit does not rely on the reasonability of other units within the organization, it may want to reserve capacity exclusively for itself. But it is still hard to see how this approach could be beneficial for the whole organization.

Total sharing is a more complicated issue because both the customer and the service provider have significant roles. Again, the customer buys a fixed capacity, probably using partial sharing and reasonable traffic estimation. The service provider may then, based perhaps on a long experience, presume that customers actually use the whole capacity reservation very rarely, if ever. Therefore, a further statistical multiplexing could be possible without too high of a risk. Yet the figures related to total sharing shown in Table 9.7 could be overoptimistic, because the quality of this service category could be of great importance for the reputation of the service provider—therefore, even a tiny risk could be too high.

9.3.4 Possible Improvements

From a customer viewpoint the main issue is how to utilize the unused capacity. If 80% or even more of the bought capacity will be unused because some applications cannot tolerate any packet losses, there could be a strong temptation to utilize the unused capacity by more tolerant applications. A technically straightforward approach is to use a BE-PHB in addition to EF-PHB. Further, as shown in Figure 9.5, it is possible to use some other PHBs to further improve the service model.

Note that in reality the situation is very convoluted. If customers can better utilize the capacity, for instance, the network operator actually needs more resources compared to the paid capacity to provide as high of quality as earlier. In the worst case, that may cause pressure to raise prices.

Figure 9.5 EF-PHB with best-effort (BE) PHB and an intermediate PHB.

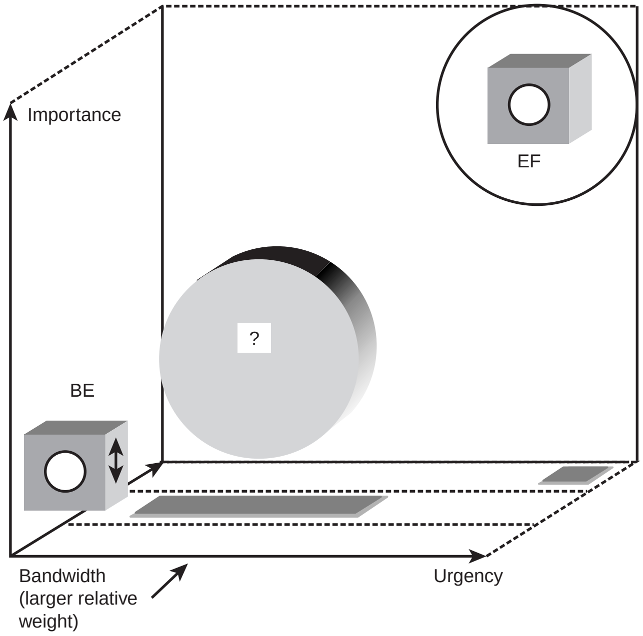

9.4 Service Differentiation with Three AF PHB Groups

One of the fundamental aims of Differentiated Services is, of course, service differentiation. The current best-effort model, albeit technically efficient, does not offer much support for advanced business models. Yet there is significant variability in the paying capacity of residential customers. This incongruity is the starting point of this implementation example. The objective is to investigate whether an AF-PHB system can provide viable mechanism for service differentiation on the Internet.

The prevalent Internet service with flat-rate pricing and best-effort service is about as simple as possible. Therefore, it is likely that the next evolutionary step cannot be based on a complex business model. Because a majority of customers likely favor simple and inexpensive service, service differentiation should mean the introduction of a couple of better service categories. One additional service is the minimum, but likely a system with three levels is a reasonable starting point. Suppose, therefore, that Fairprofit were going to build the customer service on the basis of three service grades. (Note that the term *service class* is not usable in this context because of the apparent risk of confusion with *PHB class*.)

- *Grade A*: Intended for intensive Internet use and small business as well, including the possibility of using real-time videoconferencing.
- *Grade B*: Intended for customers willing to pay more for better than best-effort service, but these customers have basically the same characteristics as ordinary users. This grade should provide clearly better service than Grade C, but with distinctly lower assured bit rates than Grade A.
- *Grade C*: Essentially the same as the current best-effort service.

In the first phase, these service grades are relevant only within the network domain administrated by Fairprofit. In the long run, however, the aim is to expand the services to other network domains as well. Moreover, Fairprofit wants flat-rate pricing to apply to all services grades regardless of the destination within the Fairprofit domain.

9.4.1 Implementation

The implementation adopted by Fairprofit is based on three AF classes. The highest AF class (AF1) is used to realize a real-time service with two importance levels. The lowest importance level (AF13) is not used because real-time characteristics are actualized by a large weight in proportion to the offered load. (That property is needed to keep delay variation small enough for the most demanding applications.)

The two highest importance levels are used to provide two levels of availability. Although the middle importance level may offer high availability, the packets with the highest importance attain preferred service during exceptional situations, such as when a primary link is broken and traffic has to be directed to an alternative route. If the link between nodes 10 and 12 gets broken, for instance, the link between 5 and 14 should be used. It is possible that the capacity of the protection link is not sufficient for all traffic demand on the middle importance level.

The second AF class (AF2) offers a non-real-time service with a better than best-effort quality. In practice, this could mean that the two highest importance levels (AF21 and AF22) provide relatively high assurance for two bit-rate levels in the same way as AF1, but without real-time characteristics. This service class may allow the use of the lowest importance level with loose or no limits of use.

The third AF class (AF3) is used to transmit best-effort packets without any additional traffic-conditioning actions in boundary nodes. In this basic model, the special characteristics of the AF-PHB class are not fully utilized because only one PHB is used in AF class 3. Figure 9.6 shows the entire AF-PHB structure.

Figure 9.6 AF-PHB structure with three AF classes.

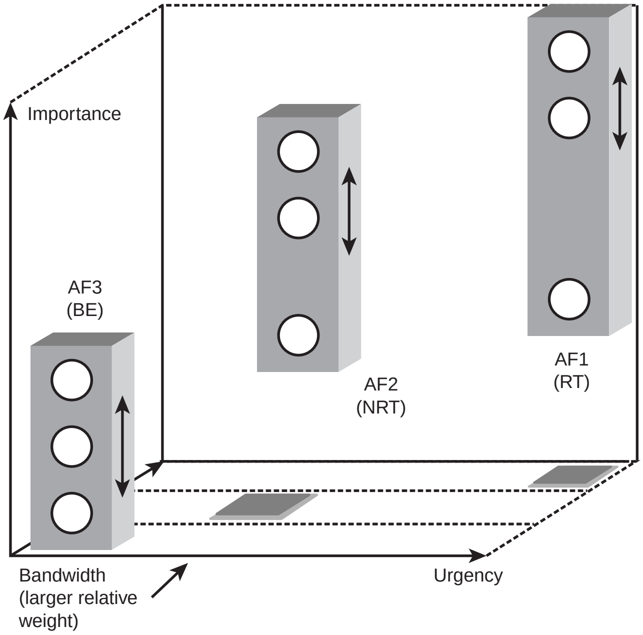

The aim of Fairprofit is to regulate the weights of each class in a way that

- The relative load level of AF class 1 is low enough to guarantee small delay variation and negligible packet-loss ratio, almost always;
- The relative load level of AF class 2 is low enough to almost always guarantee a negligible packet-loss ratio for the highest importance level, and a small packet-loss ratio for the middle importance level most of the time;
- The relative load level of AF class 3 is low enough to guarantee moderate service without significant risk of starvation for best-effort flows.

The fundamental question of the entire AF-PHB system is whether these targets are reachable by any simple management system adjusting the weights. In a real implementation, the adjusting of the weights should be automatic, without significant effort by any management personnel.

9.4.2 Traffic Model

To assess the applicability of an AF system, this section introduces a case with a large number of residential users distributed among the 14 nodes in Fairprofit's network. Each node has the following number of users:

Node 1: 100,000 users

Node 2, 3, 4, and 11: 50,000 users

Other nodes: 20,000 users

The total number of users is 480,000. This means that if every user is producing traffic with a bit rate of 100kbps, the total amount of traffic is 48Gbps. Fortunately, this number is relevant only in the access point of customers, where the possibility of statistical multiplexing is limited. In the core network, on the other hand, a significantly smaller total capacity could be sufficient.

Table 9.8 presents other traffic and service characteristics. It shows only the combined bit rate of the two highest importance levels. The middle level may represent, for instance, 80% of the total bit rate. The thresholds for real-time PHBs are systematically half of the NRT thresholds to generate incentive to use real-time service only when really needed. It is expected that Grade A and B users are not using best-effort service, although there is not necessarily any limitation of use.

For this discussion, the term *activity* means that at any point of time 5% of all possible users are active Grade C users, 0.5% are active Grade B users, and 0.05% are active Grade

A users. If a Grade A or B user is active, he or she is using real-time service with a probability of 40%. This figure may appear quite high, but notice that the duration of real-time flows are usually much longer than data flows. Moreover, real-time application may generate continuous data flows, whereas data flows often contain relatively long idle periods.

The total amount of traffic generated by a Grade A user is very high in this example. Nevertheless, that assumption could be justifiable if Grade A service is mainly used for business purposes. For instance, if all employees of a small company share one Grade A service, the traffic load could be very high during busy hour.

Table 9.8 User Profiles for Grades A, B, and C

Customer Grade	RT AF11+ AF12 kbps	NRT AF21+ kAF22 bit/s	BE (AF3x) Expected Bit Rate kbps	Activity Among All Users	RT Use	NRT Use	BE Use
Grade A	500	1000	0	0.0005	0.4	0.6	0
Grade B	100	200	0	0.005	0.4	0.6	0
Grade C	0	0	30	0.05	0	0	1

9.4.3 Network Dimensioning

The most critical task for the service provider is to dimension the network properly. The approach applied in this example is to use Formula 9.1 for three different traffic aggregates:

1. There should be enough capacity for the real-time service including both Grade A and Grade B customers. Because the load level of this service category should be low, as large a value as 20 for the parameter ψ is used. (Note that parameter ψ defines how much extra capacity is reserved for traffic variations: The larger the ψ, the smaller the probability of an overload situation.)
2. There should be enough capacity for the non-real-time service after the capacity used by the RT service is deducted from the total capacity. Because this service should provide high quality, a relatively high value, $\psi = 10$, is needed.
3. The last calculation is related to the total traffic, including all three service grades. It is reasonable to suppose that in this case a relatively low safety margin ($\psi = 5$) provides sufficient quality for best-effort service.

With the given traffic predictions, the last item produces the highest capacity requirement for all links. Table 9.9 shows only five links, because the traffic on links 1-4 and 11-12 is essentially the same as on the link 1-3. Similarly, the results on link 2-5 also are valid for links 3-6, 3-7, 4-8, 4-9, 11-13, and 11-14.

Table 9.9 Load Levels and Required Capacities (Mbps)

Link Type	RT AF11 AF12	NRT AF21 AF22	BE AF31	RT ψ=20	NRT+RT ψ=10	All ψ=5	Final Capacity
1-2	35	105	175	153	258	404	410
1-3	22	66	110	116	182	268	270
2-5	6	17	29	55	71	88	90
2-10	28	85	142	134	219	336	340
10-12	25	76	127	125	201	305	310

Next, the total link capacity of each node should be divided among the three AF classes. This example assumes that the division is equal on every link:

- RT-PHB: weight = 0.4
- NRT-PHB: weight = 0.3
- BE-PHB: weight = 0.3

The next step of evaluation is to investigate the workability of these weights. First, it is fair to assume that the weight for RT-PHB is high enough to guarantee high-quality service without noticeable packet losses. Because of the policy mechanisms of the RT-PHB, all flows using this PHB class have bit-rate limitations. Therefore, you can just subtract the expected RT load from the capacity and divide the remaining capacity evenly between NRT and BE-PHB classes because these two classes have the same weights.

The capacity within each AF class is divided between the importance levels in a manner that the highest level (for example, AF21) uses as much capacity as there is traffic demand. The middle level (for example, AF22) uses as much of the remaining capacity as there is demand. Finally, the lowest importance level (for example, AF22) can utilize the remaining capacity (if there is any).

When the network is properly dimensioned, the highest and the middle importance level of NRT flows (AF21 and AF22) cannot utilize the whole available capacity under normal circumstances. AF23 packets use the rest of the capacity available for NRT-PHB. If you suppose that all NRT users are greedy, you may assume that this AF23 capacity is divided evenly among all users regardless of the service grade. In other words, you assume that the capacity available for AF23 is divided evenly among all active users belonging to both Grade A and Grade C. An ordinary TCP implementation and basic DiffServ mechanisms probably result in that kind of division (but this issue clearly needs further studies).

The bit-rate values for NRT users are so high that all users are probably not greedy enough to use the whole available capacity. As a result, Grade C customers may in practice get a somewhat larger portion of the network capacity than what the preceding calculation indicates.

Table 9.10 Available Bit Rate per User for Grades A, B, and C

Link	Mbps	RT Mbps	share NRT=BE Mbps	NRT AF21 AF22 Mbps	NRT AF23 kbps /User	Grade A NRT kbps/ User	Grade B NRT kbps/ User	Grade C BE kbps/ s/User
1-2	410	35	188	105	214	1214	414	32
1-3	270	22	124	66	241	1241	441	34
2-5	90	6	42	17	393	1393	593	44
2-10	340	28	156	85	225	1225	425	33
10-12	310	25	142	76	236	1236	436	34

The results presented in Table 9.10 are convincing as such. It is important to note, however, that the quality of the result is largely based on the assumption that the traffic prediction during the dimensioning phase and the real traffic loads are equal. That is in reality, of course, a highly improbable situation. Table 9.11 shows a more realistic case in which the activity of customers is not known accurately. In this case, the only difference is that the number of active users at each node is supposed to be a random variable. (Actually the activity numbers have been multiplied by a random number with standard deviation of 0.3, but the total load is approximately the same as in the original case.)

On most links, the results shown in Table 9.11 do not differ much from the results in Table 9.10. It seems that flows on the lowest importance levels and on the smallest links are most vulnerable to unpredictable traffic variations. For instance, best-effort users links on link 4-8 can attain almost three times the capacity of users on link 4-9. You can also expect that the available bit rates for best-effort flows will vary widely between busy and idle hours, although this factor is not analyzed in this example.

Table 9.11 Available Bit Rate per User with Moderate Traffic Variations

Link	Mbps	RT Mbps	Share NRT=BE Mbps	NRT AF21 AF22 Mbps	NRT AF23 kbps/ User	Grade A NRT kbps/ User	Grade B NRT kbps/ User	Grade C BE kbps/ User
1-2	410	36	187	107	204	1204	404	31
1-3	270	22	124	67	235	1235	435	34

Link	Mbps	RT Mbps	Share NRT=BE Mbps	NRT AF21 AF22 Mbps	NRT AF23 kbps/User	Grade A NRT kbps/User	Grade B NRT kbps/User	Grade C BE kbps/User
1-4	270	21	124	64	261	1261	461	35
2-5	90	8	41	24	201	1201	401	31
2-10	340	29	155	87	213	1213	413	32
3-6	90	8	41	24	200	1200	400	31
3-7	90	5	43	14	587	1587	787	57
4-8	90	4	43	11	758	1758	958	68
4-9	90	10	40	30	97	1097	297	24
10-12	310	26	142	79	218	1218	418	32
11-12	270	22	124	66	237	1237	437	34
11-13	90	5	43	14	571	1571	771	56
11-14	90	7	41	22	237	1237	437	34

9.4.4 *Possible Improvements*

One approach to develop this AF service model is that the normal best-effort traffic obtains the highest importance level of AF class 3, while the other two levels are used for less than best-effort service. For instance, there could be a mechanism that marks packets with lower importance levels when it detects inappropriate behavior. These packets may encounter a very high dropping probability.

The main difficulty of this AF system is the management of weights. Although it is possible to provide three levels of differentiation, a consistent implementation could be difficult in large networks. Nonetheless, the assessment of this critical issue remains tentative until there is real experience with the management of an AF-PHB system.

9.5 *Total Service on the Basis of a DRT-PHB Group*

All the previous implementation examples were about separate issues. Yet the reality of an Internet service provider is that the same network infrastructure has to be used for all purposes. Is the right solution merely to combine the previous three approaches in the same network? That seems possible because the first implementation example required one AF class, the second one required an EF-PHB, and the last one required three AF-PHB classes.

Basically it is possible to just integrate these five PHB classes with appropriate weights. However, the management of all the weights in a way that all the differing targets are met could be a laborious effort for service providers. Therefore, the approach in this implementation example is to design a consistent framework for all services in the network. The objective is to build a network wherein one extensive service model can satisfy the various needs of different customers.

9.5.1 Implementation

The implementation of this example is based on the DRT-PHB model. It means essentially that the frame of the system is based on two PHBs classes with six PHBs. One PHB class (2) is intended for real-time applications, and the other PHB class (3) is intended for all other applications. In addition to this basic system, one PHB class with one PHB (16) is reserved for services with very high quality. It can be applied to build VPNs in the same way as EF-PHB in the second implementation example. PHB (16) needs a separate buffer and strict traffic control in boundary nodes. Another PHB (72) is reserved for network control traffic. Figure 9.7 illustrates the entire system.

Figure 9.7 DRT-PHB structure with 14 PHBs.

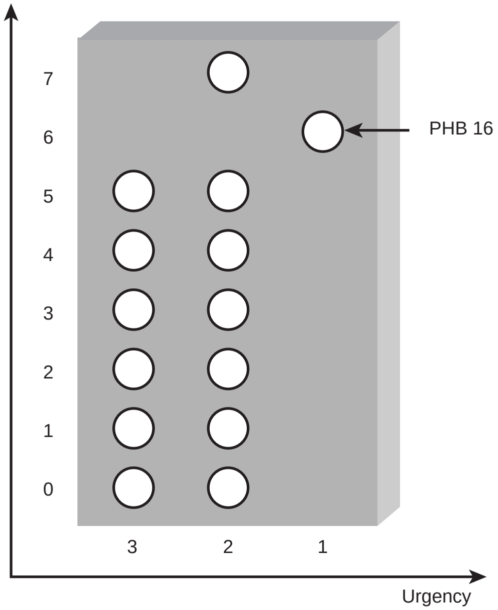

The customer service is based on the concept of *nominal bit rate* (NBR). NBR can be attached to an individual flow or aggregate stream. The primary pricing model is flat rate, which means that each customer (an individual end user or an organization) buys a permanent NBR—for instance, 50kbps for an ordinary user, or 50Mbps for a large organization.

In both cases, the boundary node measures the incoming traffic flow and compares the result with the NBR. If the measured bit rate is higher than NBR, all packets get a low PHB marking. Correspondingly, a high importance level can always be attained by sending traffic with a bit rate lower than NBR. Table 9.12 shows some examples of packet marking.

Table 9.12 Importance Level as a Function of NBR and Measured Bit Rate

NBR kbps	Measured Bit Rate, kbps	Importance Level
100	25	5
100	50	4
100	80	3
100	100	3
100	130	3
100	200	2
100	400	1
100	800	0
10000	2500	5
10000	10000	3

The same packet-marking system is applicable to both PHB class 2 and PHB class 3. There should, however, be an incentive to primarily use non real-time service. There are various possibilities to realize this kind of incentive. One simple model is that NBR is smaller for real-time service (class 2) than for normal service (class 3).

The same measuring system can also be used with PHB (16). However, the marking is entirely different. There are only two importance levels: very high (6) and immediate drop (actually less than importance level 0). The price of PHB (16) should also be consonant with the value and quality of the service, implementation especially compared to the actual price of PHB (25).

9.5.2 Traffic Model

The approach adopted here is to utilize the traffic patterns of the three previous implementation examples. Table 9.13 summarizes some key figures that define the traffic load in the network. Table 9.13 does not present the total traffic values because some users are supposed to exploit all available capacity.

The most important single figure is the NBR sum of active customers. If the NBR sum is multiplied by the average number of hops needed to transmit a packet through the network, you obtain a good estimate for the real traffic demand within the network, and thereby an estimation for the total required network capacity. In general, if that value is considerably smaller than the total network capacity, there is a good possibility for the operator to provide appropriate service.

Table 9.13 Traffic Parameters for Fourth Implementation Example

User Type	Number of Users	NBR kbps	Activity	Average NBR in Use, Mbps
University user, TCP	50000	20	0.2	200
University user, UDP	1500	20	0.2	6
Grade A users, RT	4324	500	0.022	48
Grade A users, NRT	4324	1000	0.033	144
Grade B users, RT	43243	100	0.022	96
Grade B users, NRT	43243	200	0.033	288
Grade C users	432432	30	0.056	721
Large units (VPN)	4	125000	1	500
Small units (VPN)	20	45000	1	900
In total				2903

9.5.3 Performance Evaluation

To keep the analysis simple, only the bottleneck link, 10-12, is evaluated. Because traffic figures were taken from previous implementation examples, you also can apply the link capacities from those three examples. The university reserved 20Mbps (in both directions) from link 10-12. If the partial-sharing principle with $\psi = 10$ is applied, VPNs require together 174Mbps. Finally according to the third implementation example, residential services require capacity of 340Mbps on link 10-12. Together, this means that 534Mbps implementation is supposed to be sufficient for all services.

Table 9.14 presents the results of the evaluation. The first column shows the number of active users sending traffic going through link 10-12 (all figures are for one direction). The second column shows the traffic load used by VPNs. It is expected that VPN traffic will always use the highest importance level, either 5 or 6 depending on the system adopted by the service provider. The remaining capacity, 481Mbps, is available for other customers.

For the sake of simplicity, you can expect that real-time applications send traffic with a constant bit rate regardless of the packet-loss ratio. That means that UDP users in universities always get the lowest importance level, and therefore obtain only poor service (if any).

Grade A and B customers sending real-time traffic are assumed to use real-time NBR with bit rates of 500kbps and 100kbps, respectively. When this traffic load of 26Mbps is deducted from the total capacity, the remaining 455Mbps is available for all TCP users. The last column shows the bit rate available for each customer type based on the assumption that the system can divide the total available capacity in proportion to the NBR of each user.

Because of the improved statistical multiplexing, most of the customers attain a higher bit rate than in the previous implementation examples. The only exception is Grade B users who seem to get a considerable gain by sharing the same AF class with Grade A users in the second implementation example.

Table 9.14 Available Bit Rates on Link 10-12 with a Capacity of 534Mbps

	Number of Active Users	Level 6 Mbps	Level 3 Mbps	Sum of TCP NBRs	Available Bit Rate for TCP Users kbps
Univ., TCP	1600	0	0	32	39
Univ., UDP	48	0	0	0	0
Grade A, RT	17	0	9	0	0
Grade A, NRT	25	0	0	25	1936
Grade B, RT	170	0	17	0	0
Grade B, NRT	254	0	0	51	387
Grade C, BE	4240	0	0	127	58
VPNs		53	0	0	0
In total		53	26	235	—
Remaining Mbps		481	455		

The results in Table 9.14 appear promising, although it should again be stressed that this evaluation is very limited in the sense that numerous details were ignored. For instance, the capability of the system to divide the total available capacity in proportion to NBRs should be evaluated by comprehensive simulation studies.

As a final stage of this evaluation, implementation consider a case in which link 10-12 gets broken and all traffic has to be routed through a secondary route going through the small capacity link 5-14. Table 9.15 presents a case in which the secondary route can provide only a half of the capacity of the original route. In the example system, this is no problem for the traffic with the highest importance level (VPN traffic). On the contrary, there could be significant quality differences on the lower importance levels.

Because the total capacity is less than the sum of NBRs of active customers, for example, the network cannot provide loss-free service if customers are sending traffic with their NBR. It is reasonable to suppose that because of this, real-time users of Grades A and B change to importance level 4 by reducing the bit rate to half of the original value. Still the results seem appropriate, although the situation in a real network is certainly more complex.

Table 9.15 Available Bit Rates on Link 10-12 with a Capacity of 267Mbps

	Number of Active Users	Level 5 Mbps	Level 4 Mbps	Sum of TCP NBRs	Allowed Bit Rate in Proportion to NBR kbps
Univ. TCP	1600	0	0	32	17
Univ. UDP	48	0	0	0	0
Grade A RT	17	0	4	0	0
Grade A NRT	25	0	0	25	855
Grade B RT	170	0	8	0	0
Grade B NRT	254	0	0	51	171
Grade C	4240	0	0	127	26
VPNs		53	0	0	0
In total		53	13	235	-
Remaining cap. Mbps		214	201		

9.5.4 Possible Improvements

The basic model presented here implementation seems to work appropriately in most cases. Nevertheless, there are various possibilities to modify the service model. The university may decide to prefer a more advanced system with specific NBR for each end user. That kind of system makes it possible that if someone really needs a high-quality real-time connection, a high enough NBR can be allocated for the user. Actually, the customer relationship between the service provider and university could be that the university buys a large NBR and divides that NBR among the end users.

The main advantage of the VPN model with strict traffic control and high importance marking is evident during exceptional situations, such as presented in Table 9.14. In contrast, during normal situations the model seems to waste resources, at least from the viewpoint of the organization that pays a lot of money for the network service. Therefore, rather than a typical VPN a similar model as described in the previous paragraph could be used by any organization.

If the system allows the organization to allocate the network resources (NBRs) dynamically among end users, the result could be at the same time flexible and cost efficient. There may, however, be some concerns about the robustness of the system, because in this case there is not any clear separation of the VPNs with regard to the capacity division.

Note, however, that the service model based on NBRs makes it possible to use as many service grades as the service provider wants to have, because the main tool of differentiation is NBR. The main additional feature that is probably useful for building practical service differentiation is dynamic NBR. Even in cases where flat-rate pricing is the basic model, it is possible to build a system in which each customer can acquire additional NBR based on time-dependent pricing.

The main difficulty in that approach is that the majority of the residential traffic is directed toward the customers. Consequently, a dynamic NBR approach probably requires some kind of signaling to inform the sender that the receiver is willing to pay extra NBR to get better implementation service.

Summary

This chapter evaluated the applicability of three PHB systems. Assured Forwarding seems to be a feasible approach to improve best-effort service, although the advantages are diminished if load levels are highly variable. Similarly, Expedited Forwarding seems to be a feasible approach to integrate leased-line services with other IP traffic. A DRT-PHB structure could be a practical approach to integrate differing demands into one consistent model.

With all DiffServ models, as well as with any service model, a good result requires a clear target and careful designing. What more should be said about the design and objectives of Differentiated Services? Recall the six questions asked at the end of Chapter 2, "Traffic Management Before Differentiated Services." Based on the models and analysis presented in this book, you can provide the following answers.

Q: **How can you sell a service package to ordinary customers without any technical background?**

A: If the customer service is not comprehensible, even an extremely advanced system could be worthless. Therefore, all the PHB models, such as AF-PHB group, EF-PHB, and DRT-PHB, should be hidden from ordinary end users. (The inner meaning of PHB could far too difficult to explain.)

Q: **What kind of billing system do you need to support your service model and to make it fair?**

A: A smooth evolution from the prevalent Internet model requires that flat-rate pricing be included in the total model. It is better to start with a simple flat-rate pricing scheme and then later add a more complicated pricing scheme if necessary.

Q: **Do you understand all interactions between the building blocks of services, and do they allow efficient troubleshooting?**

A: Management problems are likely if the total system consists of incompatible parts. Therefore, it is highly recommendable to plan a total system that can offer acceptable service for all purposes rather than to build separate services and mechanisms optimized for differing purposes.

Q: **How efficient is the model when used in a large network with millions of users?**

A: Something that works perfectly with a small number of users in a small network could be unsuitable for large networks. Therefore, although dynamic provision of bit rates and quality could be useful by itself, the management of all the necessary parameters could a laborious effort and prone to errors. Whenever an equal service for a majority of users is acceptable, use that model.

Q: **Is the service model robust enough to limit the effects of intentional misuse of network resources?**

A: Robust and consistent behavior is one of the key requirements of DiffServ networks. The main tool against theft of network resources is that if a flow uses more than its fair share of resources, the packets of the flow should be dropped rather than other packets.

Q: **Does the service model provide a realistic evolution path from the current best-effort network?**

A: It is unrealistic to suppose that a new infrastructure will totally replace the current best-effort network. Therefore, the treatment of best-effort traffic should be good enough to avoid starvation of best-effort flows in DiffServ networks.

Even after the considerations in this and earlier chapters, all answers tend to be somewhat vague. Various issues are still extremely difficult to assess. On a technical level, for instance, the detailed behavior of flows using TCP in case of several importance levels is still an unclear issue. Similarly, the effect of weight adjustment of several PHB classes is an extremely intricate phenomenon. Those technical issues are relatively clear, however, compared to the phenomenon of human behavior.

It is possible that the Differentiated Services model, or perhaps another new technology or business model, will change the whole picture of Internet service provision. End-user reaction to this new paradigm can only be guessed.

Fortunately, several issues probably remain the same. Some users are seeking high quality and are willing to pay more than others. They require a service with high and predictable quality most of the time (high availability of service). Quality requirement may be related to delay, packet-loss ratio, and bit rates with various combinations. Still, most users favor moderate, but inexpensive service. The promise of Differentiated Services is that it is possible to meet various demands within one network in a consistent manner.