The Assured Forwarding PHB group

- A typical application using the AF PHB is that of a company which uses the Internet to interconnect its geographically distributed sites and wants an assurance that IP packets within this intranet are forwarded with high probability as long as the aggregate traffic from each site does not exceed the subscribed information rate
- It is desirable that a site may exceed the subscribed profile with the understanding that the excess traffic is not delivered with as high probability as the traffic that is within the profile
- It is also important that the network does not reorder packets that belong to the same microflow

The Assured Forwarding PHB group

- Four AF classes are defined, where each AF class is in each DS node allocated a certain amount of forwarding resources
- IP packets that wish to use the services provided by the AF PHB group are assigned by the customer or the provider DS domain into one or more of these AF classes according to the services that the customer has subscribed to
- Within each AF class IP packets are marked with one of three possible drop precedence values
- In case of congestion, the drop precedence of a packet determines the relative importance of the packet within the AF class
- A congested DS node tries to protect packets with a lower drop precedence value from being lost by preferably discarding packets with a higher drop precedence value

The Assured Forwarding PHB group

- In a DS node, the level of forwarding assurance of an IP packet thus depends on
  - how much forwarding resources has been allocated to the AF class that the packet belongs to
  - what is the current load of the AF class
  - in case of congestion within the class, what is the drop precedence of the packet
- For example, if traffic conditioning actions at the ingress of the provider DS domain make sure that an AF class in the DS nodes is only moderately loaded by packets with the lowest drop precedence value and is not overloaded by packets with the two lowest drop precedence values, then the AF class can offer a high level of forwarding assurance for packets that are within the subscribed profile and offer up to two lower levels of forwarding assurance for the excess traffic

The Assured Forwarding PHB group

- In a DS node, the level of forwarding assurance of an IP packet thus depends on
  - how much forwarding resources has been allocated to the AF class that the packet belongs to
  - what is the current load of the AF class
  - in case of congestion within the class, what is the drop precedence of the packet
- For example, if traffic conditioning actions at the ingress of the provider DS domain make sure that an AF class in the DS nodes is only moderately loaded by packets with the lowest drop precedence value and is not overloaded by packets with the two lowest drop precedence values, then the AF class can offer a high level of forwarding assurance for packets that are within the subscribed profile and offer up to two lower levels of forwarding assurance for the excess traffic

The Assured Forwarding PHB group

- Formally, an IP packet that belongs to an AF class *i* and has drop precedence *j* is marked with the AF codepoint AF*ij*, where 1 <= *i* <= 4 and 1 <= *j* <= 3
- An AF implementation must attempt to minimize long-term congestion within each class, while allowing short-term congestion resulting from bursts
- This requires an active queue management algorithm, such as RED/RIO
- Thus, long-term congestion is managed by dropping packets, while short-term congestion (packet bursts) is controlled through buffering
- The AQM's smoothing implemented by RED/RIO effectively computes a smoothed congestion level and the dropping algorithm uses this smoothed congestion level to determine when AF packets should be discarded

The Assured Forwarding PHB group

- The AF PHB group provides delivery of IP packets in four independently forwarded AF classes
- Within each AF class, an IP packet can be assigned one of three different levels of drop precedence
- Thus, the AF PHB group included twelve AF PHBs
- The table reports the values of the DS field for the 12 available AF PHBs, as specified by RFC 2597
- A DS domain may not necessarily implement 12 AF PHBs, a suitable subset of PHBs can be selected, under the responsibility of the provider

	Class 1	Class 2	Class 3	Class 4
Low drop prec	001010	010010	011010	100010
Medium drop prec	001100	010100	011100	100100
High drop prec	001110	010110	011110	100110

The AF PHB group: Olympic services

- The AF PHB group could be used to implement, for example, the so-called *Olympic service*
- The Olympic service consists of three service classes: bronze, silver, and gold
- Packets are assigned to these three classes so that packets in the gold class experience lighter load (and thus have greater probability for timely forwarding) than packets assigned to the silver class
- The same kind of relationship exists between the silver class and the bronze class
- A Static Priority scheduler could be used to implement this relationship among the gold, silver and bronze PHBs
- If desired, packets within each class may be further separated by giving them either low, medium, or high drop precedence

The Best-Effort (default) PHB

- In the DiffServ architecture there exists a specific PHB to deliver the traditional best-Effort service
- It DS codefield is 00000000
- This is the default PHB, meaning that if a PHB is to be assigned to a packet/flow and there are not enough elemetns to decide which PHB is to be assigned, the default PHB is selected
- Best-effort service may be summarized as "I will accept your packets"
- Packets in transit may be lost, reordered, duplicated, or delayed at random
- Application traffic that uses default forwarding is expected to be "elastic" in nature
- The sender of traffic will adjust its transmission rate in response to changes in available rate, loss, or delay (TCP behavior)
- For the basic best-effort service, a single DSCP value is provided to identify the traffic, a queue to store it, and active queue management to protect the network from it and to limit delays

The Class-Selector (CS) PHB group

- The last PHB group currently specified is the Class Selector group of PHBs
- The CS PHB group serves several purposes, including the transport of certain types of multimedia applications
- Class Selector provides also support for historical codepoint definitions and PHB requirement
- The Class Selector DS field provides a limited backward compatibility with legacy practice, for example, the priority bits in the old TOS field of the IPv4 packet
- No attempt is made to maintain backward compatibility with the "DTR" or Type of Service (TOS) bits of the IPv4 TOS octet

The Class Selector Codepoints

- A specification of the packet forwarding treatments selected by the DS field values of 'xxx000|xx', or DSCP = 'xxx000' are reserved as a set of Class Selector Codepoints
- PHBs which are mapped to by these codepoints MUST satisfy the Class Selector PHB requirements in addition to preserving the Default PHB requirement on codepoint '000000‘
- We refer to a Class Selector Codepoint with a larger numerical value than another Class Selector Codepoint as having a higher relative order while a Class Selector Codepoint with a smaller numerical value than another Class Selector Codepoint is said to have a lower relative order
- PHBs selected by a Class Selector Codepoint should give packets a probability of timely forwarding that is not lower than that given to packets marked with a Class Selector codepoint of lower relative order

Compatibility with IP Precedence

- The set of codepoints 'xxx000' can be used, irrespective of bits 3-5 of the DSCP field, to yield a network that is compatible with historical IP Precedence use
- Thus, for example, codepoint '011010' would map to the same PHB as codepoint '011000', that is, when used as an old IP precedence behavior, only three bits of the DSCP are relevant
- Anyway, a Class Selector PHB does not necessarily require a static priority scheduling discipline
- Other schedulers might be used to implement a CS PHB
- This is a choice of the network administrator