

IP Quality of Service in the IETF



Kathleen Nichols
ICC 2000
June 19, 2000



This talk is not a tutorial on diffserv...

...but you can find one at:

```
ftp://ftp-eng.cisco.com/ftp/kmn-  
group/docs/hoti.pdf
```

The diffserv WG web page is at:

```
http://www.ietf.org/html.charters/d  
iffserv-charter.html
```

IntServ: the IETF's first pass at QoS




- Influenced by old-world thinking about networks
- Fundamental unit is the flow
- AS borders are invisible
- Theoretical results on treatment of flows used to specify packet treatment at each router

Diffserv: In Contrast...




- Matched to the IP architecture
- Fundamental unit is the aggregate
- Domain boundaries are modeled explicitly since they map to policy and control
- Forwarding path elements are specified based on usefulness and implementation criteria



A founding principle of the DS WG came from Dave Clark's architectural insight stated at the December 97 IETF that:

`"rules + mechanisms = services"`

A key point was that the "mechanisms" were simple primitives that could be made available in routers and composed in different ways (via the "rules") to make "services" which need to be flexible and to evolve.



The DS WG has been turning this high level architectural insight into standards and guidelines that can be used to create IP quality of service solutions, thus, in typical “standards” approach, we have evolved Clark’s simple equation to:

(rules, through traffic conditioning and classification) + (per-hop forwarding behaviors) = (an aggregate of packets with particular attributes)

which we then “simplify” by use of acronyms...



Diffserv Agenda in one slide

- Use minimal standardization to provide tools and “knobs” that
 - Allow traffic engineering within a domain
 - Allow services to be offered to customers
- Uses a bit-field in the packet (6 bits of the IPv4 TOS or IPv6 Traffic Class octet) as a codepoint to determine the packet’s forwarding treatment by indexing a table

Elements of a Diffserv Approach

- Packets are marked with codepoints (DSCPs) at the boundaries of DS domains
- Packets with the same DSCP get the same per-hop forwarding behavior (PHB) (providing aggregation and scalability)
- Per-flow state stays at network edges
- Marking can be based on microflow identification, ingress link, measured temporal characteristics of a micro- or macro-flow, etc. (Diffserv-capable equipment includes configurable traffic conditioners.)

Diffserv WG Approach

- Specify the forwarding path while letting the policy mechanisms evolve
 - analogy to IP forwarding/routing
 - this allows a wide range of uses of the basic building blocks
- Standards track for DS field and a small base set of PHBs (with limited RFC791 backward compatibility)
- Informational track for architecture doc, traffic conditioners, “per-domain behaviors”



Moving Forward in Diffserv

- PHBs have been defined that express a range of forwarding treatments (RFCs 2474, 2597, 2598)
- RFCs and WG drafts also define BA and MF classifiers, TCBs, TSW and TCM traffic conditioners for the forwarding path
- Time to specify the relationship between these in order to get particular characteristics for a traffic aggregate (with the same DSCP) across a domain



Per-Domain Behaviors (PDBs)

The expected treatment that an identifiable or target group of packets will receive from “edge to edge” of a DS domain. A particular PHB (or, if applicable, list of PHBs) and traffic conditioning requirements are associated with each PDB.


- A PDB is not a “service” but a technical specification that might be used in the creation of services (draft-ietf-diffserv-pdb-def-00 soon)
- A PDB spec is where FP and CP interact

How does this help?



- Technical attributes are defined inside the boundaries of administrative policy which avoids complicating totally technical forwarding path features.
- Modularizes end-to-end QoS, provided the PDBs aggregate in a known way
- The “next level” of putting together PDBs defined across DS domains can then be tackled. Difficulties are both technical and business issues.

Concatenating PDBs



- PDB attributes must be made “visible” at DS boundaries
- DS boundaries interior to an AS should be easier to negotiate
- PDBs that will be useful interdomain have a known behavior when merged and split

Where From Here?



- Use the toolbox, both in deployment and to develop desirable intra- and inter- domain behaviors
- This process should be fun and profitable
- There are some deployments
- Most of the research work is focusing on traffic conditioners or imaginative PHBs
- More good analysis work is needed. But that's a different talk...

SCALING THE INTERNET



Packet Design