

Inter-AS Class of Service

Implementation status and test results

Thomas Martin Knoll
 Chair of Communication Networks
 Chemnitz University of Technology
 Chemnitz, Germany
 knoll@etit.tu-chemnitz.de

I. EXTENDED ABSTRACT

Based on the presented IETF draft activities on "Cross-layer and Cross-Domain QoS signalling using BGP" at EuroView 2008, this poster presentation will document the achieved implementation status and test results on real world default free zone (DFZ) routing information.

The concept supporting Class of Service interconnection between service providers has evolved into a two fold free to join concept supporting traffic separation with consistent cross-domain packet marking and class overload prevention means.

The BGP based signalling has been implemented in the Linux Quagga routing suite and field trials with commercial routers and global routing table loads have proven its feasibility.

Measurements on transmission and storage resource usages ensure the concept's scaling capability in the view of the fast growth of DFZ routing tables.

The concept of Inter-AS CoS is applicable in IPv4 and IPv6 scenarios and should be an integral part of Future Internet strategies.

Draft-knoll-idr-qos-attribute addresses the transitive signalling of traffic class markings across layers. DSCP markings, Ethernet priority and MPLS TC bits are common examples for 6 or 3 bit supported traffic classifications in the signalling. The IP prefix originating AS announces the acceptable markings to upstream peers and respects those markings for incoming packet streams. Fig. 1 and Fig. 2 depict the two-fold concept.

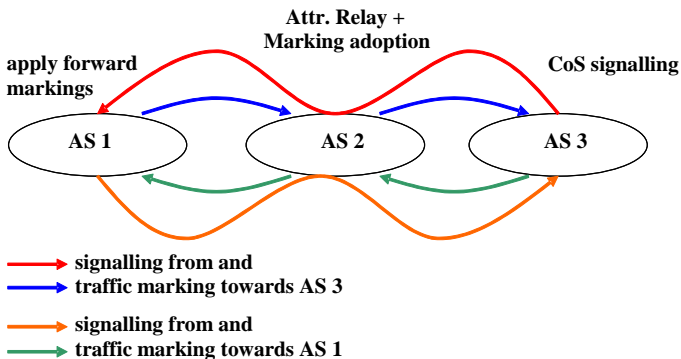


Fig. 1 Cross-Domain CoS marking

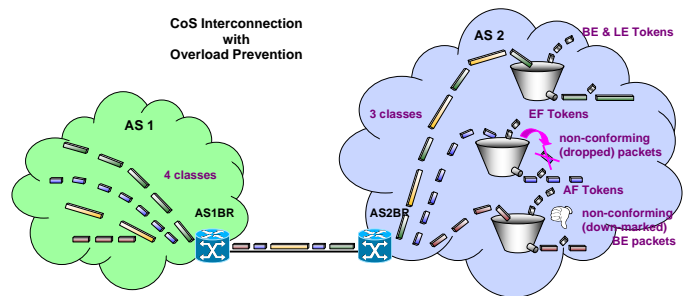


Fig. 2 CoS overload prevention

Draft-knoll-idr-cos-interconnect is the second half of the concept and addresses the likely misuse of offered higher priority classes. IP prefix advertising and relaying ASes can now signal token bucket limitation parameters for each of the supported classes in order to protect the internal associated resources and at the same time announce those ingress limitations to peering partners in a fair and square manner.

The Linux routing suite Quagga has been extended in order to support the sending and reception of those CoS signalling in its BGP module. Lab test have shown that the transitive attributes are relayed, displayed and debug recorded by commercial routers. Several measurements on real world DFZ routing tables with about 280000 routes have been performed and revealed stable router operation even under the view of thousands of inserted attributes. Each signalled class and layer requests the sending and receiving of an extended community, which results in respective RIB memory consumption. IP prefixes with identical attribute settings do not add to the overall consumption. A realistic assumption of maximum 8 supported classes on three layers with varying encodings results in less than 1MB of additional RIB storage space.

The decoding of the signalled attributes is included in Wireshark 1.2.0 and is online available at http://www.bgp-qos.org/draft-knoll/decode_attributes.php.