Cross-Layer and Cross-Domain QoS Signalling Using BGP

8th Würzburg Workshop on IP:
Joint EuroNF, ITC, and ITG Workshop on
"Visions of Future Generation Networks"
(EuroView2008)

Thomas Martin Knoll
Chemnitz University of Technology
Communication Networks
Phone **0371 531 33246**
Email **knoll@etit.tu-chemnitz.de**

Abstract

A new IETF draft is currently drawn up, which introduces a new mechanism for QoS signalling across networking layers and networking domains. It is available since the beginning of June [**I-D.knoll-idr-qos-attribute**] and shall be **presented and explained** during the EuroView symposium.
The draft uses BGP to signal QoS class markings and link selection of different layers across networking domains. It addresses the need for consistent QoS class dependent forwarding treatment of packets and frames. Routing and forwarding strategies applied on different network layers (L4, L3, L2 and L1) within a network domain are made publicly visible and adjustable. The talk will **outline the basic mechanism for the BGP signalling extension** as well as the **QoS handling strategies**, which are enabled through this mechanism. The presented draft will be **applicable to existing Internet setups** and will certainly broaden the **view on QoS handling strategies in Future Internet**.

Table of contents

Motivation

Current QoS support in the Internet

- The current "**Best Effort**" packet transport in IP networks is currently being augmented by **locally applied traffic separation** with **prioritized forwarding** together with **costly multi-parameter ingress classification**.
- Such "**quality islands**" exist **independently**, **peer with BE** traffic, run **uncoordinated QoS concepts** and might **not** even be **known globally**.
- Complex approaches exist, which aim for **guaranteed (parameterized) QoS** support for future inter-domain peerings (e.g. [MIT_CFP]).

Proposed Improvements of the new Approach

- **Provides knowledge** about the available traffic separations and encoding. **Cross-layer mapping** is a novel feature.
- Enables route selection and marking adoption without guarantees.
- Greatly improves inter-domain packet forwarding.

QoS in this approach refers to primitive traffic separation into several classes, which will experience differently prioritized forwarding behaviour in relaying nodes. Enqueueing in separate queues is thereby aspired.

Addressed Issues

Cross-Layer QoS mapping

Network type	Supported QoS classes
IP supporting DiffServ	64 (currently 21 defined)
IP supporting ITU Y.1541	6
Ethernet (IEEE 802.3)	8 (802.1q priority tag)
MPLS	8 (E-LSP) or 2^{20} (L-LSP)
ATM	4 major QoS categories
UMTS	4 major QoS categories

- IP as layer 3 and most layer 2 mechanisms support traffic class differentiation
- The number of classes and their encoding and mapping can freely be chosen by network providers.
- Diverse usage and internal QoS strategies are not necessarily visible outside a network domain
- Internal BGP (iBGP) is one choice for domain-internal QoS policy propagation.
- Increased usage of **tunnelling mechanisms** (MPLS, CE, GRE etc.) put even more pressure on **consistent inter-layer QoS coupling**.
- **Tunnels** (virtual channels) allow for QoS-based traffic engineering, which will be **regarded as Layer 1 class differentiation**.

The aim is consistent classification and a consistent class-based forwarding behaviour on all layers of an end-to-end traffic path.

Addressed Issues (cont.)

Cross-Domain QoS signalling

- Current Practice: **Best Effort only IP** traffic **peering** between ASes
- Individual agreements on class support between neighbouring ASes
- Diverse usage and internal QoS strategies are **not visible outside an AS**
- External BGP (**eBGP**) is used for **Inter-Domain** Mapping **signalling**
- **Tunnelling** of customer traffic is **preferred** for transparent transport.

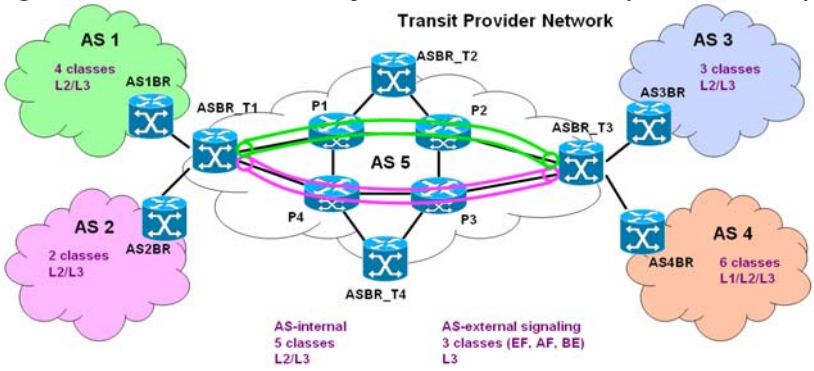

The aim is consistent classification and a consistent class-based forwarding behaviour on all layers of an end-to-end traffic path.

Definition of the QoS Marking Attribute

Ext. Community Attribute

The new QoS Marking Attribute is encoded as a **BGP Extended Community Attribute** [RFC4360]. It is therefore a **transitive optional** BGP attribute with **Type Code 16**.

The Type Value has been assigned to 0x00 [IANA_EC].

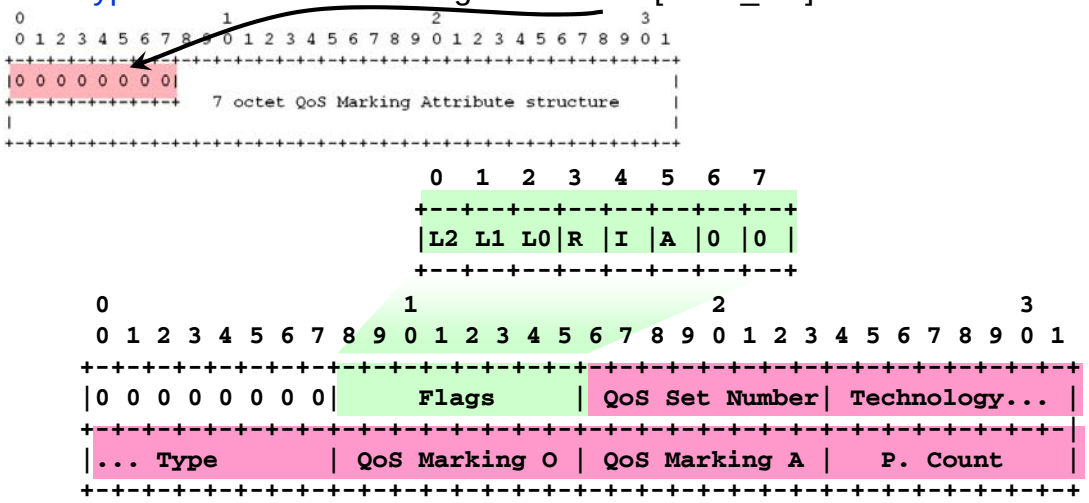

BGP update message

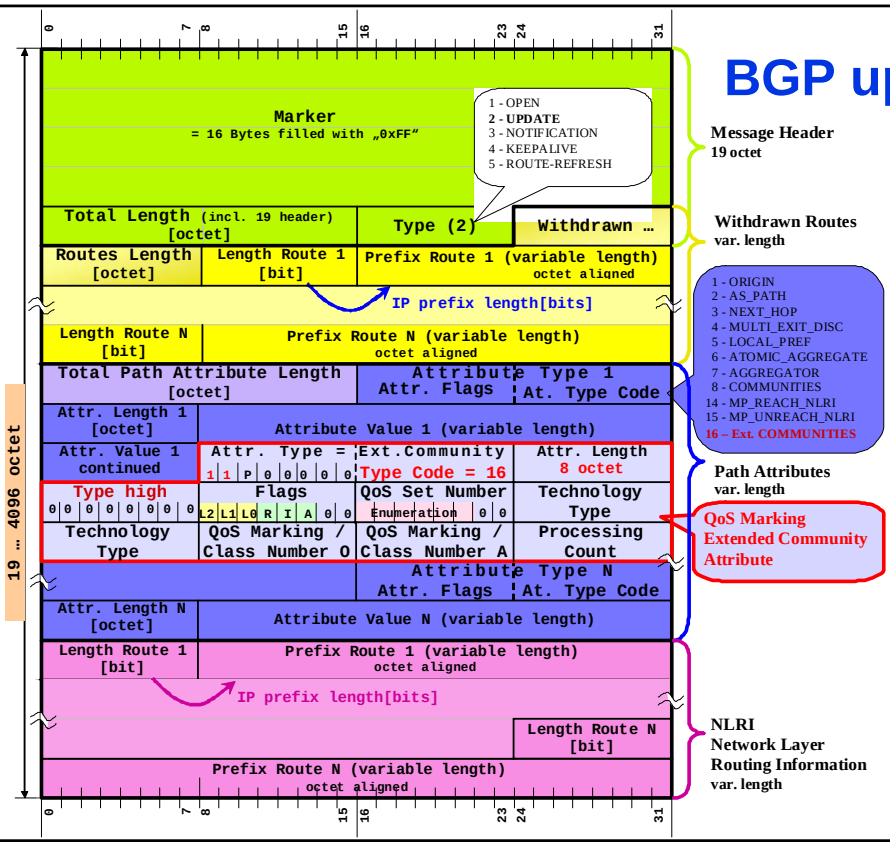

Selected Mechanisms of the Draft

Optional transitive Attribute

- Smooth integration and transparent transport across ignoring ASes
- Fixed fields guarantee unchanged values / other fields for local adaptation

QoS Set – Concept of "linked" together attributes

- Several QoS Attributes will be included, which are virtually grouped together
- Grouping not fixed to technology or DSCP etc.

Technology Type

- Lack of common enumeration of different layer technologies → L0..2 selection
- Simplified approach in next version

Processing Count

- Detection of non-cooperative ASes (Count vs. diff. AS numbers in AS_PATH)
- Route selection based on 'I' flag and P. Count possible

Ideas for future QoS Designs

General comments

- Distinction between direct peering and transit peering (avoid remarking)
  → favour tunnelled transport
- Define a general Technology Type enumeration for cross-protocol (service) consistent numbering
- L1 priority -> encompass QoS path/media selection for seamless interworking with optical and radio networks
- High need for a consistent Class of Service concept

Ideas for future QoS Designs (cont.)

Class Set Definition

1. Best solution: fixed standard including metering, enforcement and allocation
   e.g. using ITU parameters [Y.1541]
2. Free choice + class signalling + recommendations
   e.g. using “Configuration Guidelines for DiffServ Service Classes”, [RFC4594]
3. Free choice without signalling (confidential status) → not wanted
4. No Class Set support → not wanted

Class Mapping / Encoding Mapping

1. Best solution: fixed cross-layer standard including encoding and mapping
2. Free choice + class encoding signalling -> DSCP as anchor point
   (eases tunnelling and provides “inferred” QoS treatment
3. Free choice without signalling (confidential status) → not wanted
4. No cross-layer Class Set support → not wanted

Summary

- The proposed approach enables a general QoS based forwarding which allows for informed routing and marking decisions.
  It is optimized for ease of deployment and adopted to the current poor inter-domain forwarding model.

- Future designs should aim for a consistent and widely standardized QoS framework, which encompasses cross-layer and cross-domain traffic class handling from L1 to at least L3 as generally offered QoS treatment.

- More sophisticated QoS concepts are not prohibited and will always exist, which results in future “better quality islands”.

Sources

[IANA_EC] IANA, „Border Gateway Protocol (BGP) Data Collection Standard Communities“, online available http://www.iana.org/assignments/bgp-extended-communities

[I-D.knoll-idr-qos-attribute] Knoll, T., "BGP Extended Community Attribute for QoS Marking", draft-knoll-idr-qos-attribute-00 (work in progress), June 2008.

[MIT_CFP] Amante, S., Bitar, N., Bjorkman, N., and others, "Inter-provider Quality of Service - White paper draft 1.1",November 2006, http://cfp.mit.edu/docs/interprovider-qos-nov2006.pdf

[RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.

[RFC4360] Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", RFC 4360, February 2006.

[RFC4594] Babiarz, J., Chan, K., and F. Baker, "Configuration Guidelines for DiffServ Service Classes", RFC 4594, August 2006.

[Y.1541] ITU-T, “Network performance objectives for IP-based services”, Y.1541, February 2006