A Concept of Inter-AS Priority Signaling using BGP Attributes

NETWORKS 2008
13th International
Telecommunications Network Strategy
and Planning Symposium
September 28 – October 2, 2008

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

Outline

1. Motivation
2. Addressed Issues
3. Definition of the QoS Attribute
4. Selected Mechanisms of the Specification
5. Remarks and Use Cases
6. Summary

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**.
  (markings are ignored, reset and re-classified)
- 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 Concept

- **Provides knowledge** about the available traffic separations and encodings. **Cross-layer mapping** & **transitive Cross-domain signalling** 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-Layer QoS mapping (cont.)

- cross-domain tunnelling of customer traffic
  →consistent inter-layer QoS coupling
  →transparent transport

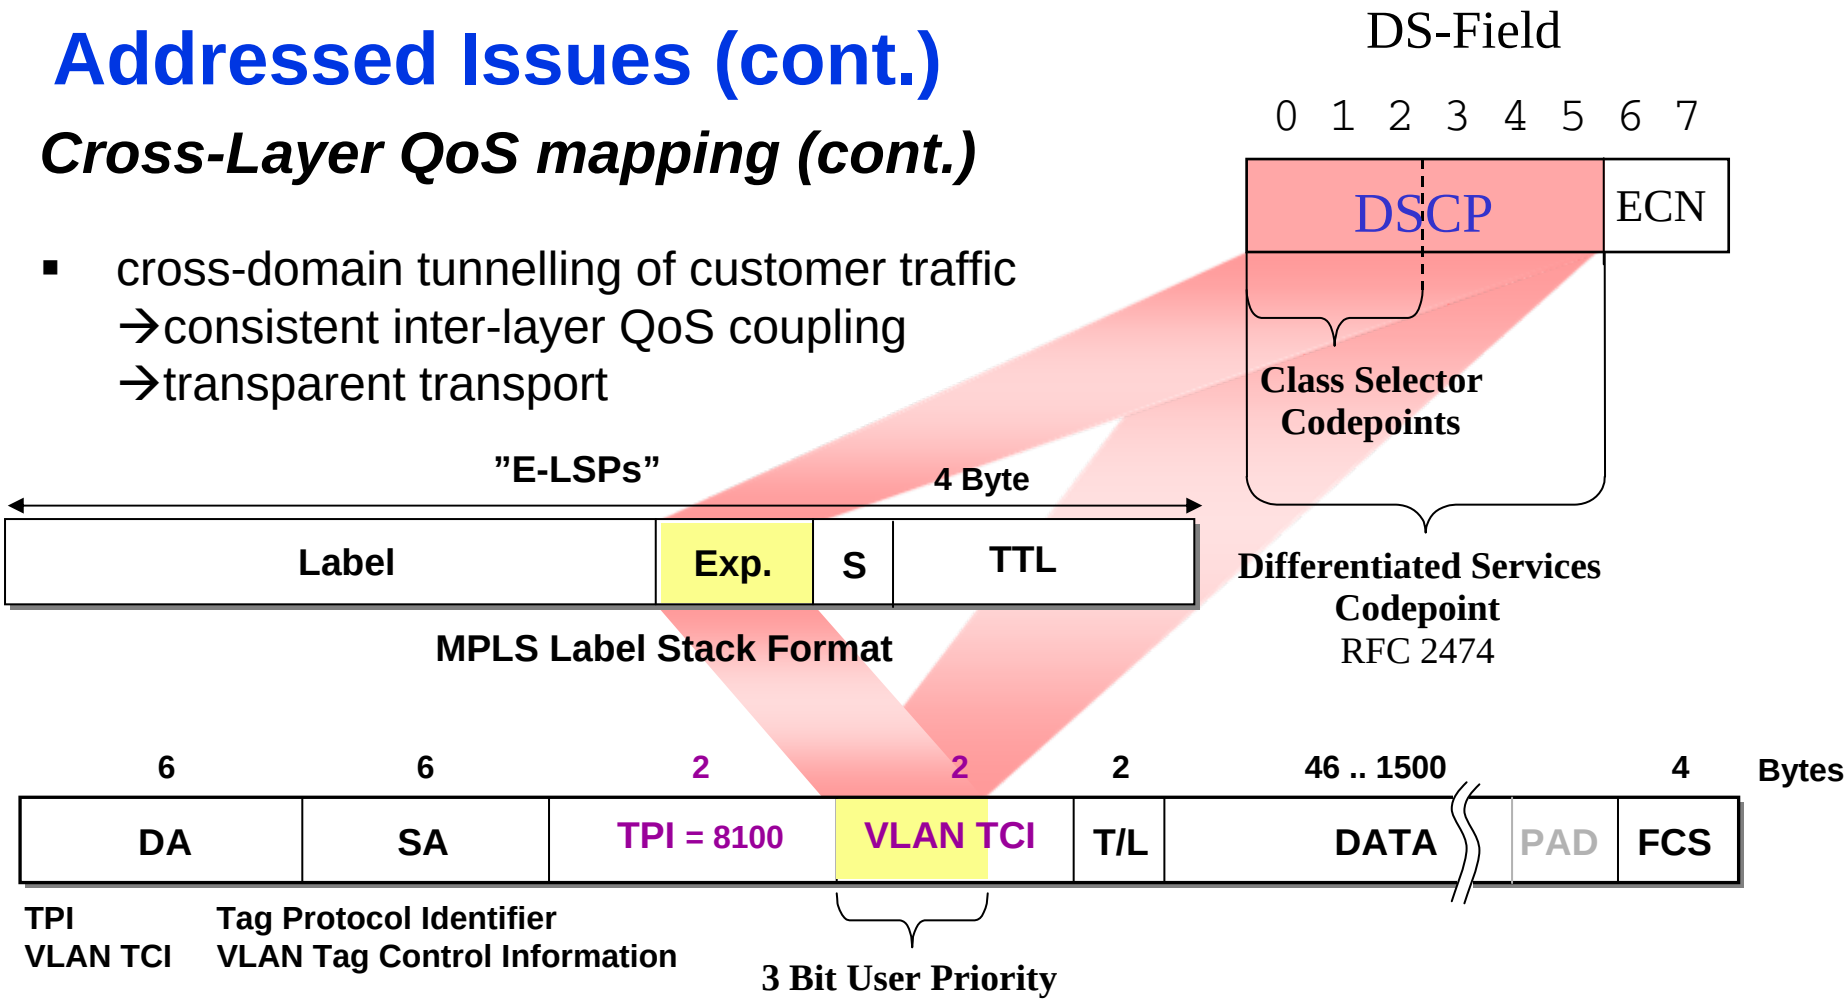

The aim is consistent classification and a consistent class-based forwarding behaviour on all layers of an end-to-end traffic path.

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.

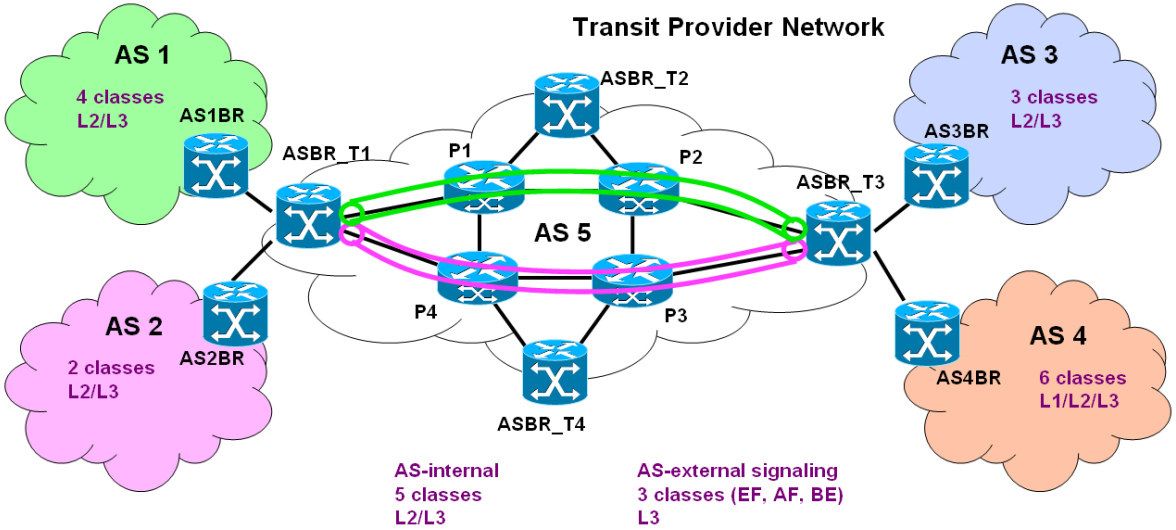

The aim is consistent classification and a consistent class-based forwarding behaviour on all layers of an end-to-end traffic path.

eBGP peering between neighbouring ASes

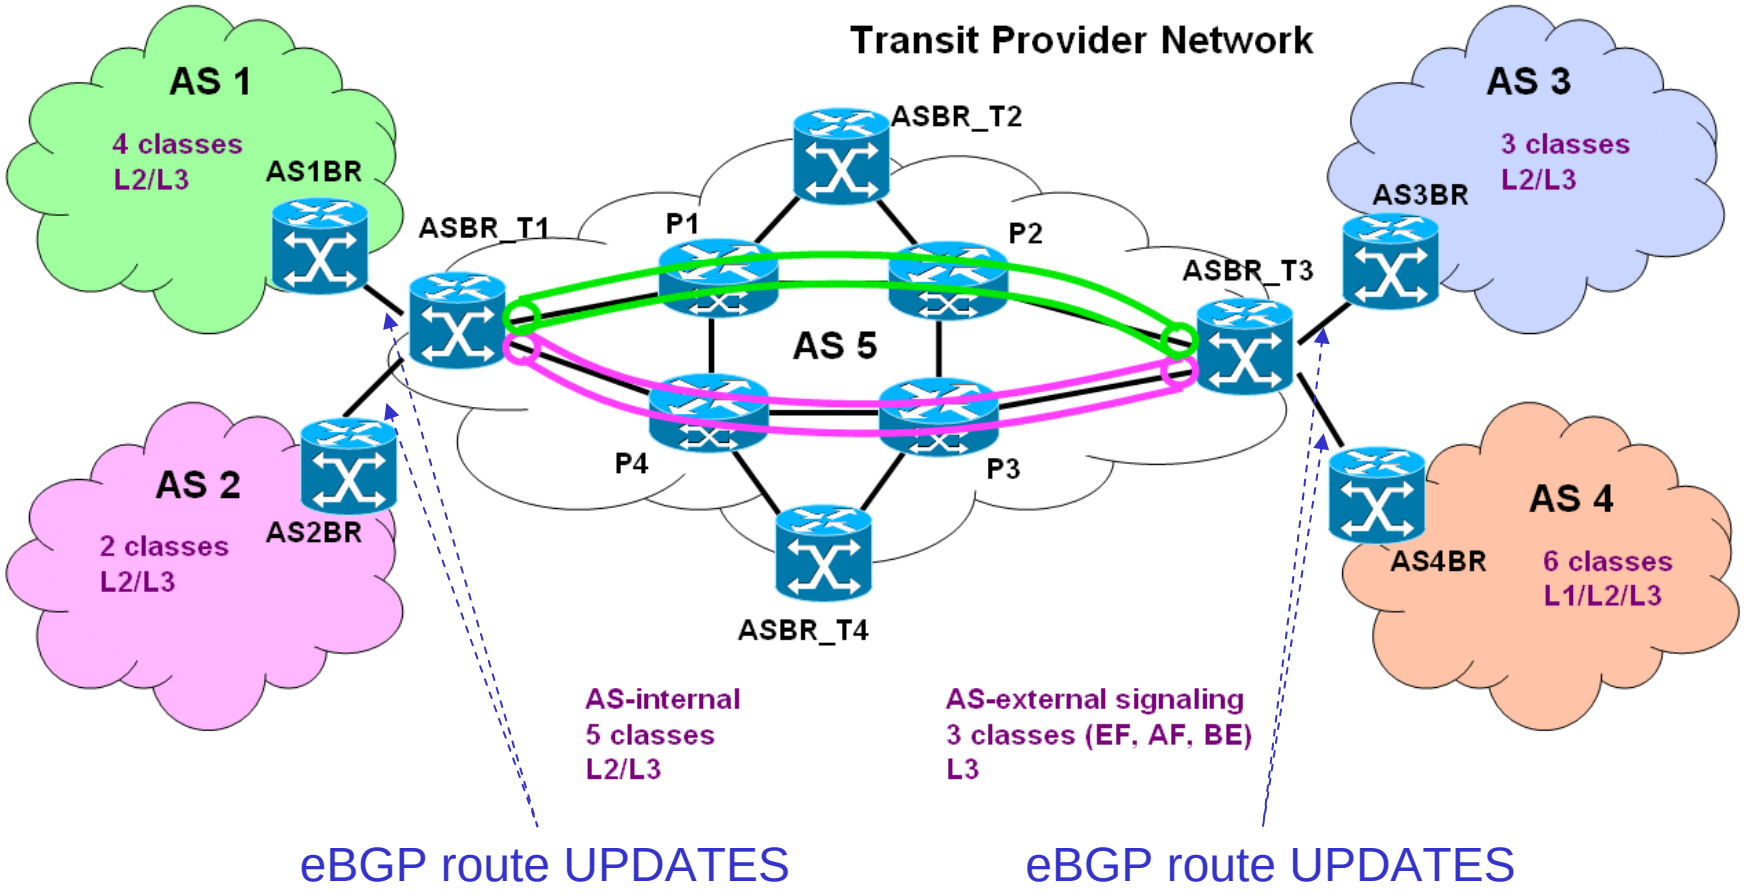

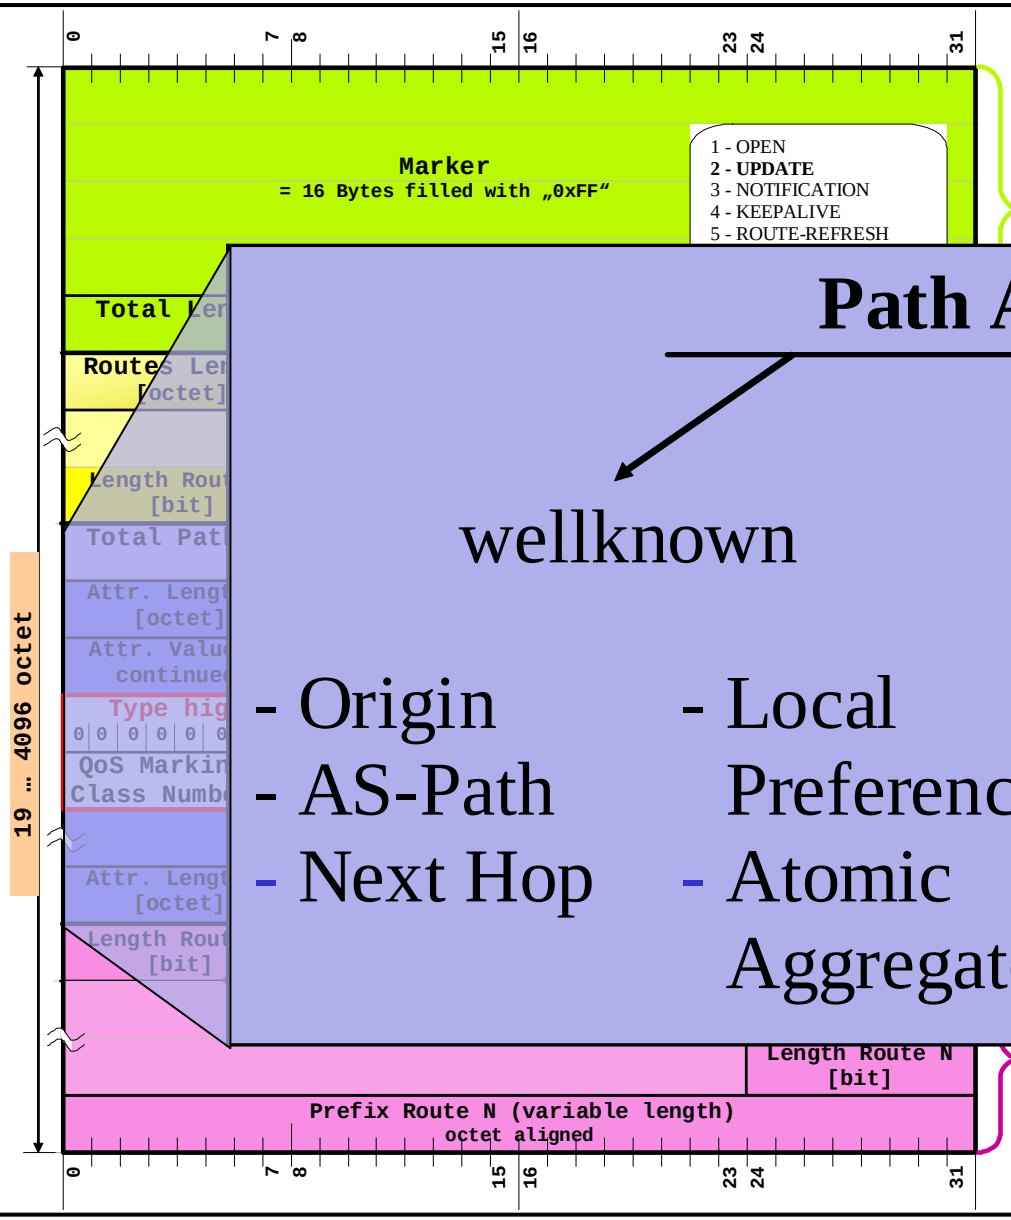

BGP update message

Path Attributes

wellknown		optional
- Origin	- Local Preference	- Aggregator
- AS-Path	- Atomic Aggregate	- MED
- Next Hop		- Community
		- Ext. Comm.

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].

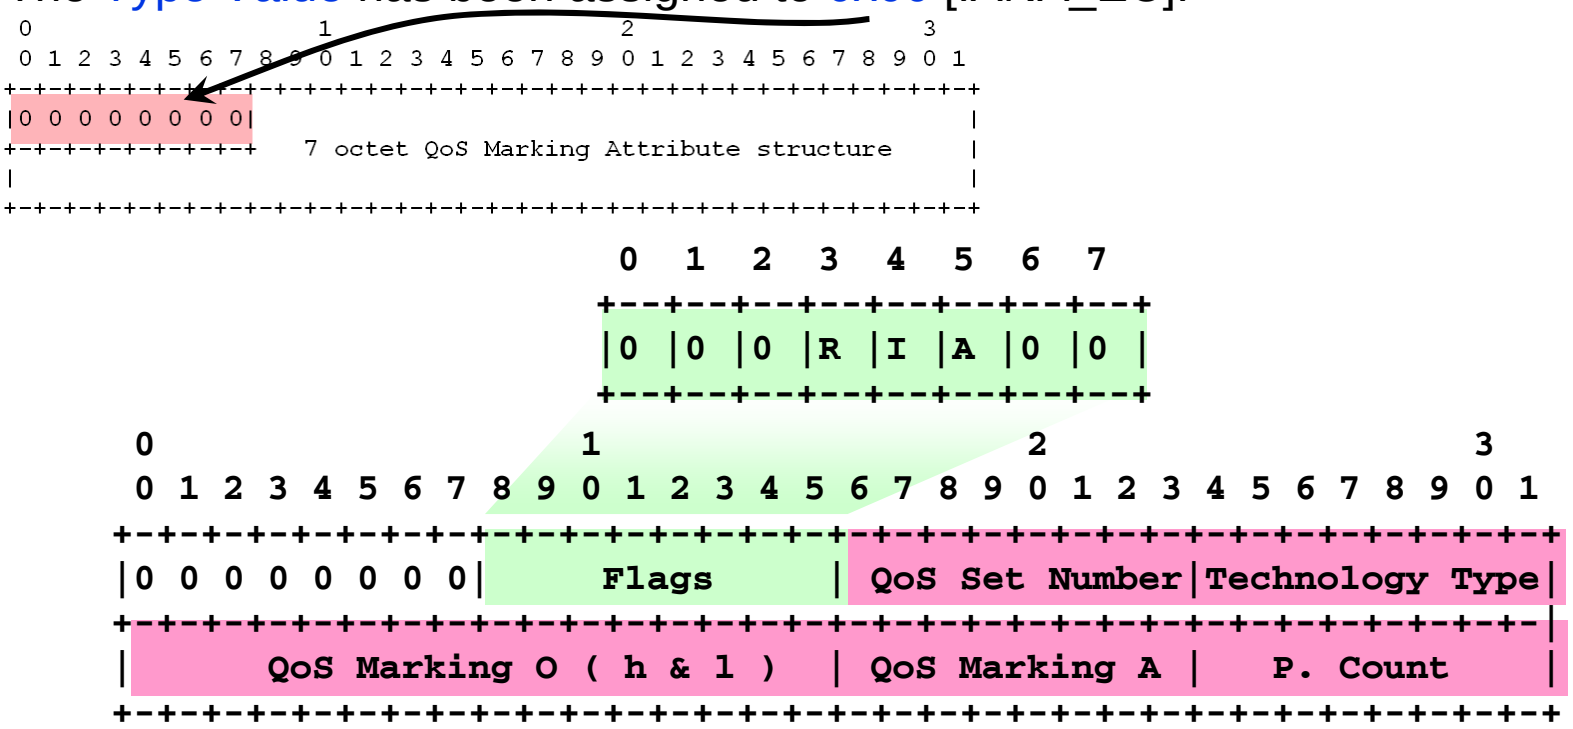

Selected Mechanisms of the Concept

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
  → own enumeration list

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
- Additional usage of the attribute's 'P'-flag

Remarks and Use Cases

General comments

- Distinction between direct peering and transit peering (avoid remarking)
  → favour tunnelled transport for transit traffic
- Define a general Technology Type enumeration for cross-protocol (service) consistent numbering -> started
- L1 priority -> encompass QoS path/media selection for seamless interworking with optical and radio networks
- Usage of the Class Set information and the processing count analysis for the best path selection process -> ongoing debate about process changes and multi-path inter-AS peerings
- Usage of the Class Set information and the processing count analysis for PCE calculations
- High need for a consistent Class of Service concept

Remarks and Use Cases (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.
- The concept aims for a **consistent and widely adopted QoS approximation**, which encompasses cross-layer and cross-domain traffic class handling from L1 to at least L3 as generally offered QoS treatment.
- The concept **incorporates a confidentiality option** that allows operators the distinction between an secluded internal and the advertised external Class Set.
- **More sophisticated QoS concepts** are **not prohibited** and will always exist, which results in future “better quality islands/path”.

Outlook & Questions

Not part of today's presentation, but complementary

BGP Class of Service Interconnection

http://tools.ietf.org/html/draft-knoll-idr-cos-interconnect-00

Fancy a general 3 class Internet ?