IP網におけるトラヒック及び輻輳制御
－ITU-T勧告Y.iptcの紹介と今後の展開－

米田　進

日本テレコム(株) 情報通信研究所〒104-0032 東京都中央区八丁堀2-9-1
E-mail: susumu.yoneda@japan-telecom.co.jp

あらまし　インターネットの急速な発展に伴いＩＰ網におけるトラヒック及び輻輳制御に関する研究開発及び実用化が進展している。遅ればせながら国際標準の場（ITU-T）においてもこの分野における取組みを開始し、この程勧告案が承認され、勧告化手続きを進めている。本稿は、その勧告案の内容並びに他のデファクト標準との差異を紹介すると伴に、今後の展開について述べる。

キーワード　IP、トラヒック制御、輻輳制御、IPパケット転送能力

Traffic Control and Congestion Control in IP based Networks
－The Introduction of ITU-T Recommendation Y.iptc－

Susumu Yoneda

Japan Telecom Information and Communication Labs. 2-9-1 Hatchobori, Chuo-ku, Tokyo, 104-0032 Japan
E-mail: susumu.yoneda@japan-telecom.co.jp

Abstract　A rapid expansion of The Internet stimulates the R&D activities as well as the implementations of traffic control and congestion control technologies in IP based networks. At international standardizations such as ITU-T, the study of traffic and congestion control in IP based networks is started, and in January of this year the draft recommendations was approved by the study group. This paper introduces the concepts, definitions and procedures addressed in this recommendation. These approaches of traffic and congestion controls are compared with the existing Per Hop Behaviors in IETF. And some future study items are discussed.

Keyword　IP, Traffic Control, Congestion Control, IP Packet Transfer Capability

1. Introduction

This document introduces the ITU-T Recommendation, Y.iptc, which was technically approved in the past January. This recommendation states that the primary role of traffic control and congestion control procedures is to protect the IP network and the traffic entering the network in order to achieve the IP network performance objectives and QOS commitments. Therefore, the traffic control and congestion control addressed here are the required procedures for committing the QOS level. However, it will be impossible for all nodes of the Internet to implement the procedures. It may be practical for a limited portion of IP based network to implement them so that the IP performance objectives and commitments could be met within this limited portion of the network.

In IP based networks, congestion is defined as a state of network elements (e.g., routers, switches) in which the network is not able to meet the network performance objectives and the negotiated QOS commitments for the established IP flow. An IP flow refers to the set of IP packets categorized into a certain type of classifications. The classification is based on the information in IP packet fields such as source/destination addresses and TOS/DS.

For some IP flows, the network commits to meet QOS objectives, assuming the user generated IP packet flow conforms to a traffic contract. For IP flows without a traffic contract between the user and the network, the network may provide a best effort service. In order to define a traffic contract, we have to define IP transfer capabilities, which characterize IP flows.

In this document, we address traffic parameters and descriptors in Section 2, and Section 3 presents three different IP transfer capabilities to meet three different types of traffic contracts. Section 4 addresses differences from IETF Intserv/Diffserv.

2. Traffic Parameters and descriptors

IP flow and IP subflow

An IP flow at a given interface is defined as the occurrence at that interface of the set of IP packets which match a given classification. An IP flow may consist of packets from a single application session, or it may be an aggregation comprising the combined traffic from a number of application sessions. When a classification may be subdivided into different sub-classifications (separate or overlapping), different IP subflows may be recognized in the corresponding IP flow.

Traffic Parameter

A traffic parameter describes one aspect of a flow. It may be qualitative or quantitative. A traffic parameter may for example describe the peak bit rate, the peak packet rate, the average bit rate, the average or maximum packet size, the (average of maximum) burst length of a flow, etc.

Traffic Descriptor

A traffic descriptor is the set of traffic parameters that is used to capture the traffic characteristics of an IP flow at a given standardized interface as part of the traffic contract.

Traffic Contract

For a given IP flow, the selected IP transfer capability, the traffic descriptor at a given interface and the QOS class (see Recommendation Y.1541 [ITU1541]) define the traffic contract at that interface.

3. IP transfer capabilities

An IP transfer capability is a set of network capabilities provided by IP based networks to transfer IP packets. For each IP transfer capability, the service model, traffic descriptor, conformance definition and any QOS commitments are defined. An IP transfer capability is supported by a set of traffic control and congestion control functions.

In order to offer multiple classes of QOS to multiple applications and to optimize the usage of network resources, IP based networks should be capable of providing multiple transfer capabilities.

Three IP transfer capabilities are defined:

- Dedicated Bandwidth (DBW) IP transfer capability
- Statistical Bandwidth (SBW) IP transfer capability
- Best-effort (BE) IP transfer capability

This set of IP transfer capabilities is based on current IP service models and this set may be extended in the future.

3.1 Dedicated Bandwidth (DBW) transfer capability

Description

The Dedicated Bandwidth (DBW) transfer capability is intended to support applications with stringent delay requirements. It aims to support the guaranteed and timely delivery of IP packets along the end-to-end path of the network.

The DBW transfer capability strives for compatibility with the Guaranteed Service [RFC2212] and the end-to-end services based on the Expedited Forwarding per-hop behavior [RFC2598].

Service model

The DBW IPTC can be used by applications that characterize the traffic with a single token bucket.

The commitment made by the network is that the negotiated IP QOS is assured to all IP packets when all packets are conforming to the conformance tests. The DBW user should expect that (possibly all) non-conforming packets be discarded by the network.

The DBW capability can be associated with specified loss commitments (IP Loss Ratio, IPLR) and specified delay commitments (IP Transfer Delay, IPTD and IP Delay Variation, IPDV), (see Recommendation Y.1541).

The network does not fragment packets. In addition, the network commits to attempt, as long as possible (e.g. until there is a need to reroute the flow), to maintain packet sequence integrity.

Traffic descriptor

The Traffic Descriptor consists of the following parameters:

- The peak rate and peak bucket size.
- The maximum allowed packet size.

Conformance definition

An IP packet is conforming if the arrival conforms to the following two parts:

- the arrival is conforming to the Generic Byte Rate Algorithm (GBRA) [ITU371];
- the actual packet length does not exceed the maximum allowed packet size.

The GBRA is updated for conforming packets only.

QOS commitments

The DBW capability may be associated with specified loss commitments and specified delay commitments.

If all packets are conforming, the QOS commitments apply to all IP packets. The DBW user should expect that (possibly all) non-conforming packets be discarded by the network. If not all packets are conforming, the network may choose to commit QOS to some of the packets, for example to a volume of packets that is conforming.

3.2 Statistical Bandwidth (SBW) transfer capability

Description

The Statistical Bandwidth (SBW) transfer capability is intended to support applications, which do not have stringent delay requirements. It aims to support the guaranteed delivery of IP packets along the end-to-end path of the network.

The SBW transfer capability strives for compatibility with the Controlled Load Network Element Service [RFC2211] and the end-to-end services based on the Assured Forwarding per-hop behavior [RFC2597].

Service model

The SBW transfer capability provides a specified sustainable rate for non-real time applications with limited burst duration with the expectation that traffic in excess of GBRA will be delivered within the limits of available resources.

The following two examples describe the commitment the SBW user will receive:

- If the user sends conforming packets at a constant rate that is less than or equal to the sustainable rate, then the commitment is that all these packets are delivered across the network, corresponding to the associated QOS class.
- If the user has not sent packets for a long time and the user sends conforming packets in a burst with a duration that does not exceed the bound set by GBRA [ITU371], then the commitments is that all these packets are delivered across the network, corresponding to the associated QOS class.

The SBW capability also allows the user to send conforming packets in excess of the GBRA, but traffic that exceeds this bound will only be delivered within the limits of available resources.

The SBW capability may be associated with a specified packet loss commitment.

The network does not fragment packets. In addition, the network commits to attempt, as long as possible (e.g. until there is a need to reroute the flow), to maintain packet sequence integrity.

Traffic descriptor

The traffic descriptor consists in:

- The peak rate and the peak bucket size;
- The sustainable rate and the sustainable token bucket size;
- The maximum allowed packet size.

Conformance definition

An IP packet is conforming if the arrival conforms to the following three parts:

- the arrival is conforming to the peak GBRA;
- the arrival is conforming to the sustainable GBRA;
- the actual packet length does not exceed the maximum packet size.

The GBRAs are updated in coordinated mode for conforming packets only.

QOS commitments

The SBW capability may be associated with specified loss commitments.

If all packets are conforming, the QOS commitments apply to all packets. Otherwise, the QOS commitments apply to a number of bytes in conforming packets. Non-conforming traffic will be delivered within the limits of available resources.

3.3 Best effort (BE) transfer capability

Description

The best effort IP transfer capability is intended to support applications which do not have stringent loss or delay requirements.

Service model

The service model for the best effort (BE) IPTC requires that available resources be used for forwarding packets of best effort flows. Even though there is no QOS commitment specified, the expectation is that packets be delivered provided that sufficient resources are available.

Traffic descriptor

- The maximum allowed packet size.

Conformance definition

Conformance to the maximum packet size is required.

QOS commitments

There is no absolute QOS requirement for this transfer capability.

4. Differences from IETF Intserv/Diffserv

This document describes two IP transfer capabilities (i.e. DBW and SBW) that strive for compatibility with service descriptions and specifications generated by IETF IntServ and DiffServ groups. The intention is that network elements and networks conforming to these specifications, are likely able or adaptable to support the IP transfer capabilities specified in this document.

This section lists and motivates the most important differences between the IP transfer capabilities and the IETF counterparts.

4.1 Dedicated BandWidth IPTC associated with a suitable QOS class

The description for the DBW IP Transfer Capability lists that the DBW capability strives for compatibility with the Guaranteed Service (GS) [RFC2212] and the end-to-end services based on the Expedited Forwarding per-hop behavior (EF) [RFC2598].

Because ITU-T uses a combination of an IPTC and a QOS class, the following assumes that a QOS class with specified loss commitments and specified delay (variation) commitments is selected for a given flow. A flow characterized by the DBW capability in combination with such a QOS class is in the sequel referred to as 'DBW' . Similarly, a flow which uses the GS or a flow using the EF PHB is referred to as 'GS' and 'EF' respectively.

The common property between DBW, GS and EF is that the network commits to transport IP packets with a specified capacity [byte/s] with a low delay and low delay variation.

The following differences are noted.

- The DBW capability allows the network provider to discard (possibly all) non-conforming packets, whereas in GS and EF non-conforming packets are expected to be carried on a best-effort basis. It is expected that the ability to restrict the influx of DBW packets to the agreed rate and characterised by GBRA, greatly simplifies the traffic control in the network (e.g. priority queuing) and minimises queuing delays for all conforming traffic. It is further expected that most applications using DBW (e.g. streaming media) are able to specify suitable parameters and to generate traffic within these bounds (conforming packets).
- The association with a QoS class allows the end-to-end commitments to be a priori (before agreeing on a traffic contract) known to the user. GS allows (e.g. via the RSVP protocol [RFC2210]) to inform the user about the maximum queuing delay expected to be experienced on the particular flow path. This information is provided a posteriori (after the reservation has been made). In addition, it has been shown that on a path with many hops the accumulation of per-hop maximum queuing delay values leads to much larger values than the 'maximum delay' (or a suitable quantile thereof) experienced on that path. Consequently, the accumulated information is of little or no practical value to the user. EF does not specify an upper bound to the delay (variation).

4.2 Statistical BandWidth IPTC associated with a suitable QOS class

The description for the SBW IP transfer capability lists that the SBW capability strives for compatibility with the Controlled-Load Network Element Service (CL) [RFC2211] and the end-to-end services based on the Assured Forwarding per-hop behavior (AF) [RFC2597].

Because ITU-T uses a combination of an IPTC and a QOS class, the following assumes that a QOS class with specified loss commitments and possibly un-specified delay (variation) commitments is selected for a given flow. A flow characterized by the SBW capability in combination with such a QoS class is in the sequel referred to as the 'SBW'. Similarly, a flow that uses the

CL or a flow using the AF PHB/PSC is referred to as 'CL' and 'AF' respectively.

The common property between SBW, CL and AF is that the network commits to transport IP packets with at least a specified capacity [byte/s] and allows additional (excess) traffic to be transported within the limits of the available network resources.

The following differences are noted.

- The SBW capability and AF is intended to support applications that do not have stringent delay requirements. CL does not commit to quantified delay properties but aims at properties 'under unloaded conditions', i.e. a delay variation that is small compared to the minimum delay (e.g. propagation delay). This property is expected to be difficult to realise in combination with the ability to efficiently support excess traffic beyond and in addition to the guaranteed rate.
- Currently, the SBW capability distinguishes conforming and non-conforming packets. AF distinguishes three so-called 'drop precedence', yet is required to support at minimum two different drop probabilities.

5. Conclusion

This paper introduces the concepts, definitions and procedures addressed in the recommendation Y.iptc. These concepts and procedures are fundamentally different from the PHBs defined in IETF. One of such differences is that IPTC can provide a known QOS level to users before they start the sessions, but PHBs can address the spontaneous QOS level when users want to start the sessions.

IPTCs are recently defined, and no implementations can be seen at this point. However, they use the same technologies and techniques implemented for ATM networks.

When considering the exact network design and QOS commitments, we may need to have techniques such as IPTCs so that QOS levels can be controlled to meet user's requests. Nowadays an IP based network is a critical part of many of businesses. In this environment, the industry may need to address technologies and techniques to exactly control QOS for the requested IP flows.

Reference

[ITU371] ITU-T Recommendation I.371 (2000): Traffic Control and Congestion Control in B-ISDN

[ITU1540] ITU-T Recommendation Y.1540 (1999): Internet Protocol Communication Service – IP packet transfer and availability performance parameters

[ITU1541] Draft ITU-T Recommendation Y.1541 (2002): Internet Protocol Communication Service – IP Performance and Availability Objectives and Allocations

[RFC2210] IETF RFC 2210 (1997): The Use of RSVP with IETF Integrated Services

[RFC2211] IETF RFC 2211(1997): Specification of the Controlled-Load Network Element Service

[RFC2212] IETF RFC 2212 (1997): Specification of Guaranteed Quality of Services

[RFC2698] IETF RFC 2698 (1999): A Two Rate Three Color Marker

[RFC2475] IETF RFC 2475 (1998): An Architecture for Differentiated Services

[RFC2597] IETF RFC 2597 (1999): Assured Forwarding PHB Group

[RFC2598] IETF RFC 2598 (1999): An Expedited Forwarding PHB