Zeitschrift: Comtec: Informations- und Telekommunikationstechnologie =

information and telecommunication technology

Herausgeber: Swisscom Band: 78 (2000)

Heft: 6

Artikel: Intelligent bandwith management and end-to-end quality of service

Autor: Gysel, Peter

DOI: https://doi.org/10.5169/seals-876452

Nutzungsbedingungen

Die ETH-Bibliothek ist die Anbieterin der digitalisierten Zeitschriften auf E-Periodica. Sie besitzt keine Urheberrechte an den Zeitschriften und ist nicht verantwortlich für deren Inhalte. Die Rechte liegen in der Regel bei den Herausgebern beziehungsweise den externen Rechteinhabern. Das Veröffentlichen von Bildern in Print- und Online-Publikationen sowie auf Social Media-Kanälen oder Webseiten ist nur mit vorheriger Genehmigung der Rechteinhaber erlaubt. Mehr erfahren

Conditions d'utilisation

L'ETH Library est le fournisseur des revues numérisées. Elle ne détient aucun droit d'auteur sur les revues et n'est pas responsable de leur contenu. En règle générale, les droits sont détenus par les éditeurs ou les détenteurs de droits externes. La reproduction d'images dans des publications imprimées ou en ligne ainsi que sur des canaux de médias sociaux ou des sites web n'est autorisée qu'avec l'accord préalable des détenteurs des droits. En savoir plus

Terms of use

The ETH Library is the provider of the digitised journals. It does not own any copyrights to the journals and is not responsible for their content. The rights usually lie with the publishers or the external rights holders. Publishing images in print and online publications, as well as on social media channels or websites, is only permitted with the prior consent of the rights holders. Find out more

Download PDF: 15.11.2025

ETH-Bibliothek Zürich, E-Periodica, https://www.e-periodica.ch

Exploration Programmes: Corporate Technology Explores Future Telecommunications

Intelligent Bandwidth Management and End-to-End Quality of Service

Customer "A" buys a data service over an IP network. When using business-critical applications for the first time, unexpected problems appear. Voice quality over the IP network turns out to be insufficient. People using applications such as the inter-enterprise software packages from SAP encounter very long response times. Although the network operator assures that the performance of the IP backbone corresponds to the service level agreement, customer "A" experiences a poor performance of important applications and believes he is not getting what he has paid for. What has happened? And what can the operator do to improve the customer satisfaction?

The Programme "Network Architectures & Technologies" explores the emerging IP functionalities supporting fixed and mobile services, and identifies key solutions to engineer and plan next generation packet based networks. The objective is to achieve cost reduction in network investment and operation and to transfer new network capabilities into revenue generating converged network services. With its Exploration Programmes, Corporate Technology is exploring telecommunication technologies and new service possibilities with a long-term view of 2–5 years. Further, the expertise built up in the course of this activity enables active support of business innovation projects.

hen a network operator offers a data service over an IP network, e.g. LAN Interconnect, he usually guarantees a minimal throughput and a maximum delay through the network from the ingress point to the corresponding egress point.

PETER GYSEL, BERNE

The example "LAN Interconnect" is depicted in figure 1. The Quality of Service (QoS) is defined on OSI-layer 2 and may be fixed in a Service Level Agreement (SLA). This means that the QoS is guaranteed for the sum of all traffic, including junk traffic as well as real-time traffic. Today, there is usually no QoS defined on the application layer. Hence, the end-to-end quality of an individual application encountered by the customer (QoS on layer 7) can be quite different from that expected from the SLA.

this article – also in the IP data backbone of the carrier. On the access line, however, bandwidth is limited. Non-critical or undesired, bandwidth-intensive applications may cause traffic congestion in the access line. If a customer has an access line with one Mbps throughput, one single IP-TV session may cause overflow in the downstream traffic. All other applications that need high QoS, e.g. voice or video conferencing running over UDP (User Datagram Protocol) or business critical applications running over TCP (Transport Control Protocol), e.g. SAP, will be degraded. This situation is depicted in fig. 2a). This example shows that QoS on layer 2 cannot guarantee a sufficient end-to-end quality of the applications. Three approaches may be considered to resolve the problem:

- (1) Increasing the bandwidth of the access line
- (2) Introducing QoS in the access network

- (3) Introducing bandwidth management at the customer premises The first approach results in higher cost for the customer and money will be spent for undesired traffic. Moreover, traffic congestion can still happen and bandwidth critical traffic is still not really protected. The second approach, "Integrated Services" with RSVP signalling [1], attempts to provide per-flow QoS assurances with dynamic resource reservation. This approach does not scale in a backbone environment; it would require large infrastructure investments and additional effort in network management. In this article we concentrate on the third approach with bandwidth management on the application layer at the customer premises. This approach has the advantage that the network can be used as it is. No changes are needed. Only a policyenforcer function has to be introduced at the customer premises. This is indicated in figure 1 with a dashed line.

The first question to be discussed will be how to allocate bandwidth for a distinct application. The second issue concerns the question which applications are allowed to occupy bandwidth, i.e. defining policies and administrating them.

Three products have been tested in our IP lab, one for the needs of SMEs (Small and Medium Enterprises) and two highend products covering high-performance demands. Special attention has been put on user-friendliness, traffic

Results show that a better end to end QoS on the application layer can be achieved by intelligent bandwidth management at the customer premises.

This article first describes how the quality of services like "LAN Interconnect" or "Internet Access" can be improved by intelligent bandwidth management with a policy enforcer at the customer premises. Then the main features of different scale products are presented and the results of tests are summarised. Finally, the management of large networks with many policy enforcers is considered.

Methodology

In a configuration as shown in figure 1 generally plenty of bandwidth is available in the LAN and – as is assumed in




Fig. 1. Typical configuration of a LAN-Interconnect service. The LAN of the head-quarters with a server farm is connected to the LAN of one or more subsidiaries. The dashed line indicates where a policy enforcer can be introduced. ER: Edge Router, CPE: Customer Premises Equipment.




Fig. 2. Typical traffic distribution a) without and b) with policy enforcement at the customer site.

monitoring and reporting, functionality (achieved throughput for different applications, response time for TCP, lost data for UDP), and scalability of management and administration of large networks.

Key Features of Policy Enforcers

Solution for SMEs

For SMEs, the task mainly consists in defining a few traffic flows which have to be prioritised. ColtSoho (Ascom) integrates the functionalities of a hub, a router with traffic prioritisation and xDSL transmission in a single piece of equipment (fig. 3). This results in very interesting prices. The definition of the prioritised traffic can be made e.g. according to the layer-4-protocol (UDP or TCP), the IP source address, RTP (Real-time Transport Protocol) payload types or the packet size. Within the two queues (priority queue, default queue) a "first in first out"-mechanism is used.

It was found that real time traffic such as voice over IP, video conferencing etc. can be efficiently protected from undesired aggressive traffic. In fact, no UDP packets from voice applications were lost. The number of applications which can be treated separately, however, is limited in that equipment since there are only two queues.

High Performance Solutions

Products like Packet Shaper from Packeteer Inc. and NetEnforcer from Allot Communications offer a wider range of functionalities for intelligent bandwidth management and can handle a lot of different queues. They are placed between the LAN and the customer router (dashed line in fig. 1).

First, traffic is observed and monitored. Traffic is then divided in either predefined or user-defined traffic classes. The classification criteria can be the port number of the application, the IP source or destination address, IP subnets, etc. The current use of bandwidth for the different flows as well as the accumulated value over different periods of time can be monitored and reported.

Then different policies such as minimum or maximum bit rate, priorities, never admit or discard can be defined and enforced for individual or aggregate flows. For TCP traffic this can be realised by retaining packets and acknowledgements; for UDP traffic a scheduling mechanism is needed. Furthermore the policy enforcer equipment can write and read the Differentiated Service Code Point (DSCP). Thus the policy enforcement in the access is compatible to the concept of Differentiated Services [3] for the backbone. Business critical or real-time applications are given an appropriate policy and thus

get the bandwidth they need (fig. 2). If there are no critical applications the bandwidth is left to other applications. Further interesting features of policy enforcers concern reporting. E.g., for TCP traffic the round trip time of an individual flow can be evaluated and compared with an upper limit defined in an SLA. Sometimes, it may be helpful to divide the total delay into a network component and a server component. Another feature is the accumulation of the total amount of traffic of a department which allows automated billing for user groups. The tests have shown significant differences in the functionality of the products and the implemented features. It turned out that the reservation of a certain bandwidth for a given traffic class does not always work reliably. To slow down an application seems easier than reserving a fixed bandwidth. Although the user interfaces look quite comfortable, care must be taken when defining appropriate policies. The definitions must sometimes be very precise, otherwise bad performance may be encountered.

Conclusions

The user of IP services perceives the endto-end quality of his application. This quality can differ a lot from the quality which would be available on layer 2 between the two concerned customer routers. Therefore, intelligent management of the limited bandwidth in the access line is important.

Bandwidth management at the customer premises can absolutely protect high-priority traffic from loss, even under heavy overload conditions. It may allow the customer to get better performance for




Fig. 3. Colt Soho as an example of an integrated equipment with the feature of policy enforcement.

References

- [1] R. Braden, L. Zhang, S. Berson, S. Herzog, and S. Jamin, "ReSerVation Protocol (RSVP) Version 1 Functional Specification", RFC2205, Sept. 1997.
- [2] D. Black et al., "An Architecture for Differentiated Services," IETF RFC 2475, Dec. 1998.
- [3] R. Rajan et al., "A Policy Framework for Integrated and Differentiated Services in the Internet", IEEE Network, vol. 13, no. 5, 1999.
- [4] M. Wahl, T. Howes, and S. Kille, "Leightweight Directory Access Protocol (v3)," RFC 2251.
- [5] J. Boyle et al., "The COPS (Common Open Policy Service) Protocol," Internet draft, Aug. 1999.

specific applications without the need of upgrading the access line. Policy-based bandwidth management thus optimises the use of the resources. From the network operator's point of view, policy-based networking allows SLAs on the application layer rather than the network layer.

The tests of several products have shown that the user interfaces of policy enforcers are quite easy to use. Some suppliers use advanced concepts for supporting DiffServ. However, the tests have also shown that definition of policies can be tricky and implementation should be left to a specialist.

Outlook

Future customers will need bandwidthon-demand for selected applications. For example, a customer requests a highbandwidth video session during specific times of the month. Bandwidth should only be allocated during these periods. A network operator will have to manage an enormous amount of information on a wide range of users, applications and resources. He therefore needs an administration tool that brokers information stored in static and dynamic databases. This information must be translated into policies and into commands which can be understood by diverse policy enforcement equipment [3] (fig. 4). The administration tool must keep track of the committed bandwidth so that the behaviour of the network remains predictable. Fur-




Fig. 4. Administration tool for large networks with many policy enforcers.

ther it can collect data on the network usage and updates accounting databases.

To allow end-to-end QoS for connections over two, or more, network domains,

Abbreviations

ATM Asynchronous Transfer Mode COPS Common Open Policy Service LDAP Lightweight Directory Access Protocol POP Point of Presence Quality of Service QoS RSVP Resource Reservation Protocol SLA Service Level Agreement Small and medium enterprises SME TCP Transport Control Protocol User Datagram Protocol UDP xDSL DSL (Digital Subscriber Line)

the administration tool must interface to bandwidth brokers. Bandwidth brokers from different carriers automatically negotiate SLAs. All the interfaces of the administration tool have to be based on standard protocols such as LDAP [4], COPS [5] or SNMP (Simple Network Management Protocol). This enables a network operator to deliver new services very rapidly.


Peter Gysel is with "Hochschule für Technik und Architektur Biel" which is part of the "Berner Fachhochschule".

During winter 1999/2000 he was with Swisscom Corporate Technology working in the project "Policy Based Networking" of the Exploration Programme "Network Architectures & Technologies".

Summary

technology

Die Qualität von Datendiensten über IP-Netze, wie beispielsweise LAN Interconnect oder Internet Access, kann durch intelligentes Bandbreitenmanagement für die Anschlussleitung erheblich verbessert werden. Dabei wird zwischen der Dienstqualität auf Schicht 2 des OSI-Modells und der End-zu-End Qualität von Anwendungen (Schicht 7) unterschieden. Ein so genannter Policy Enforcer, der zwischen dem LAN und dem Router des Kunden eingefügt wird, teilt verschiedenen Anwendungen die vorgesehene Bandbreite zu. Es zeigt sich, dass kritische Anwendungen auch unter schweren Überlastbedingungen vollständig vor aggressiven, bandbreiten-intensiven Verkehrsflüssen geschützt werden können.





Wer uns heute für Informatik und Kommunikation kontaktiert, profitiert schon morgen davon.

SOHARD AG - Generalunternehmen für

- Digital Audio Broadcast Solutions
- Globale Informations-Systeme wie Postphone, Bankphone, Fahrgast, Parkplatz
- Flottenmanagement-Systeme f
 ür Transportunternehmungen, Rettungs- und Pannendienste
- Oracle based Solutions
- Mobile Datenverarbeitung für Aussendienst, Service, Verkauf
- Internet, Intranet, E-Commerce
- · Service, Support, Sicherheit



Software/Hardware Engineering Galgenfeldweg 18, CH-3000 Bern 32 Tel. 031 33 99 888, Fax 031 33 99 800 E-Mail: sohard@sohard.ch Internet: www.sohard.ch

ISO 9001 Reg.-Nr. 10909-02