Resilience-Differentiated QoS – Extensions to RSVP and DiffServ to Signal End-to-End IP Resilience Requirements

Achim Autenrieth [(1)]*, Andreas Kirstädter [(2)]

(1) Munich University of Technology
Institute of Communication Networks
Arcisstr. 21
80290 Munich, Germany
Phone: +49 89 289-23518
Fax: +49 89 289-62318
E-mail: Autenrieth@ei.tum.de

(2) Siemens AG
Corporate Technology, Information and Communication
Otto-Hahn-Ring 6
81730 Munich, Germany
Phone: +49 89 636-47484
Fax: +49 89 636-51115
E-mail: andreas.kirstaedter@mchp.siemens.de

Abstract

Based on the growing commercial importance of the Internet network resilience is becoming a key design issue for future IP-based networks. In this paper an extension to existing Quality of Service (QoS) architectures is presented which integrates the signaling of resilience requirements with the traditional QoS signaling. We refer to this extended QoS model as Resilience-Differentiated QoS (RD-QoS). The applications signal their resilience requirements in addition to their QoS requirements to network edge. The network takes the resilience requirements into consideration for the resource management and traffic handling. At the border of MPLS domains, the resilience requirements can then be directly mapped to the appropriate MPLS recovery options. This approach allows an integrated end-to-end provisioning of resilience and QoS in an IP-based network employing MPLS.

Keywords

IP and MPLS, QoS Aspects, Network Architecture

1 Introduction

With the tremendous growth of the Internet, additional services with new requirements were developed. Real-time services like video-conferencing or IP telephony have a connection-oriented character and require high Quality of Service (QoS) with hard delay, delay jitter and bandwidth constraints. The Internet also became an e-commerce platform with a fast growing commercial importance. Network outages due to link or node failures result in direct loss of revenue and, moreover, loss of reputation. Therefore, e-commerce companies and business customers of the Internet require a "non-stop" Internet availability of up to 99.999% similar to present TDM networks.

Quality of Service architectures

To support the Quality of Service requirements of real-time, connection-oriented services, two QoS models were defined by the IETF: the Integrated Services (IntServ) architecture with RSVP as a signaling protocol, and the Differentiated Services (DiffServ) architecture. In the IntServ model, the QoS requirements of the services are signaled on a per-flow basis through the network and the required network resources are reserved using RSVP. The DiffServ model provides QoS to aggregated traffic on a per-hop basis. At the QoS domain boundary the traffic is classified into service classes, and the service classes are then conditioned at each router according to their negotiated service level requirements.

Even though the reliability of a service is an important attribute of the service quality, no resilience attribute is currently defined for the QoS architectures. Network survivability is treated independently of the QoS architectures.

Network Survivability

In [1] the network survivability is identified as a key requirement of traffic engineered networks. Several survivability options are defined and guidelines are given to support these requirements in MPLS-based networks.

It is generally possible to reach the survivability requirements of IP services with resilience mechanisms present at different protocol layers, e.g. SONET/SDH, ATM, WDM and MPLS [2]. Moreover, these resilience mechanisms may even be in operation in multiple layers at the same time. While recovery at lower layers generally has advantages in the time scale of the recovery operation, recovery at the IP or MPLS layer allows a better resource efficiency and finer recovery granularity. A resilience differentiated approach could protect only those traffic flows requiring a high level of service availability. This results in a more cost-effective network design and traffic engineering.

Therefore it is reasonable for an ISP to provide the required network survivability using only resilience mechanisms in the IP layer. That way, also the network operation and management complexity could be reduced, since all traffic engineering aspects (including resilience) are managed in the IP layer only. ISPs can offer unprotected and protected services (the latter at higher cost) with a single administrative platform, including user authentication and billing. This is a major advantage since it reduces the operational cost of the network and increases service flexibility. Depending on the amount of money a customer is willing to pay he receives a customized level of resilience. Customers who accept lower network resilience may be offered low-cost network services. Customers demanding high network resilience tolerate higher charges.

An open issue for an ISP is, however, how to identify services with high resilience requirements involving fast recovery mechanisms. This leads to the problem how to establish a service with high end-to-end survivability. Existing QoS architectures so far don't allow the signaling of resilience requirements. Current IP resilience strategies allow the provisioning or resilience only by (slow) network management functions. In this paper an extension of the currently discussed QoS architectures is proposed, which allows an integrated handling of QoS and resilience requirements. The extended QoS architecture that can differentiate the resilience requirements of IP services will be referred to as *'Resilience-Differentiated QoS' (RD-QoS) architecture* [3]. The objective of the proposed architecture is the end-to-end provisioning of service resilience over edge and core IP networks.

This paper is organized as follows. In section 2 the basic RD-QoS architecture is introduced and a classification of services into resilience classes with corresponding recovery options is proposed. Section 3 discusses the extensions to the DiffServ architecture and to RSVP to support RD-QoS. In section 4 the integration of the RD-QoS with MPLS is presented to realize an end-to-end IP resilience provisioning. In the final section the benfits and requirements of RD-QoS are presented.

2 Resilience-Differentiated QoS

Since the resilience requirements and the classical QoS requirements of IP services are orthogonal to each other, an extended quality-of-service definition was proposed in [4]: the combination of the commonly discussed quality-of-service in terms of bandwidth and delay together with the resilience requirements of the application. The proposed way to signal the resilience requirements is to include the corresponding signaling into the QoS signaling between the application and the network. Corresponding to the different QoS approaches (IntServ, DiffServ) this could either be done on a per flow or on a per packet basis.

RD-QoS architecture

The RD-QoS architecture extends the existing QoS architectures to support differentiated resilience requirements of IP services. The resilience requirements are included in the quality-of-service signaling between the application and the network. Depending on the QoS architecture used, the signaling may be along a full end-to-end route or between the application and the network boundary. In either case the signaling includes resilience attributes identifying the resilience requirements of the service. Packets belonging to a certain resilience class are marked accordingly at the network boundary. Depending on the QoS architecture the marking may be done using the TOS-byte (DiffServ Code Points in the IP header) or using an explicit label (MPLS) or by referring to certain flow descriptions (IntServ, RSVP).

Additionally, the network must take care that the required QoS level can be maintained in case of a network failure with a minimum of service outage time. This requires a careful bandwidth and resource management which reserves enough spare resources to allow a service continuity for a given set of expected failures, e.g. all possible single link failures in the network domain. The bandwidth management may either reserve dedicated resources on two physically disjoint paths through the network, or keep a pool of spare resources, which can be shared by multiple services in the event of failures.

Under normal conditions (i.e. without any network failure present), traffic is handled by the QoS architecture according to the negotiated service level agreement without the RD-QoS extension.

In the event of a failure however, the traffic conditioning takes the resilience requirement of the service class into consideration. Packets of low-priority pre-emptible services with no resilience requirements may be discarded to free network resources for services with resilience requirements. Depending on the negotiated level of resilience, the queuing and dropping precedence of these services may be modified.

Service classification and resilience schemes

This paper proposes a set of four resilience classes (see Table 1) primarily distinguished by their recovery time requirements.

Table 1: Proposed service classes and corresponding resilience options

Service Class	RC1	RC2	RC3	RC4
Resilience requirements	High	Medium	Low	None
Recovery time	10-100 ms	100ms - 1s	1s - 10s	n.a.
Resilience scheme	Protection	Restoration	Rerouting	Pre-emption
Recovery path setup	pre-established	on-demand immediate	on-demand delayed	none
Resource allocation	per-reserved	on-demand (assured)	on-demand (if available)	none
QoS after recovery	equivalent	may be tempo-rarily reduced	may have reduced QoS	none

- Resilience Class 1: High resilience requirements

  The resilience scheme used for services with high resilience requirements is protection switching. Both 1+1 and 1:1 protection is possible. For a 1+1 protection, packets must be forwarded on a working and an alternative path simultaneously. In case of a failure on the working path, the downstream side simply selects packets from the alternative path. In case of 1:1 protection the packets are forwarded on a predefined alternative path only in case of a network failure. The protection resources may be used for low-priority, pre-emptible traffic as long as no failures are present in the working path. This requires a recovery signaling to handle uni-directional failures.

- Resilience Class 2: Medium resilience requirements

  For medium resilience requirement, restoration techniques (or fast rerouting) may be used where the recovery path is setup after a failure detection (similar to [5], section 2.1.1). In this case spare resources are inherently shared for the protection of different working paths. On service setup, the resource management has to assure that enough spare resources are available for a given set of expected failures. In case of a network failure, packets are forwarded after a fast rerouting and reservation of spare resources

- Resilience Class 3: Low resilience requirements

  For services with low resilience requirements, recovery resources are not considered during traffic engineering processes (neither exclusively nor shared). In case of a failure, packets may be forwarded after a rerouting and reservation phase, if enough resources are available. This implies, that the services may experience reduced QoS after the recovery.

- Resilience Class 4: No resilience requirements

  In case of a network failure in the administrative domain, packets with no resilience requirements may be discarded/dropped. This may happen even if the traffic is not directly affected by the network failure but rather by a re-routing of other traffic having higher resilience requirements. This corresponds to low-priority, pre-emptible traffic in telecommunication networks.

The resilience classes define the basic resilience behavior of the service. For more efficient resource management, additional resilience attributes may be defined. These attribute could specify, if the service tolerates a reduced Quality of Service in the event of a network failure. The drawback of these additional resilience attributes is that the signaling and resource mangement is more complex.

In the next section the required functions and extensions of QoS architectures to support the different resilience classes are discussed.

3 Application to existing QoS architectures

The proposed way for signaling the resilience requirements of IP services is to include the corresponding signaling into the quality-of-service signaling between the application and the network. With the DiffServ architecture this could be done on a per packet basis while with RSVP this would be done on a per flow basis.

Extensions to RSVP / RSVP-TE

In the IntServ architecture [6] resources are reserved by the RSVP [7] for every QoS flow on every router along a path between sender and receiver. While IntServ is not widely used due to its scalability problems, RSVP evolved to a versatile signaling and reservation protocol. Several traffic engineering extensions are proposed for RSVP (RSVP-TE) to allow Constraint Based Routing (CBR).

In the RD-QoS architecture, the RSVP signaling must be extended to reserve spare resources in the network to provide survivability against network failures. The RSVP message formats are extended in the sense that the end user's terminal is able to signal a resilience requirement to the network in addition to the classical QoS requirements like bandwidth and delay (jitter).

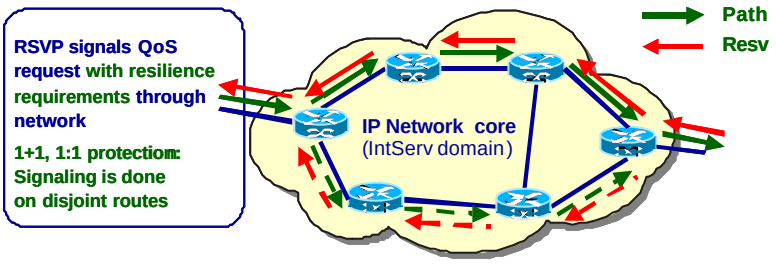

Figure 1: RD-QoS signaling with RSVP

The proposed way to do this is to include the resilience requirement in the Rspec [8] of RSVP. The three IntServ classes – Guaranteed, Controlled Load and Best-Effort – are combined with a two-bit resilience attribute identifying the resilience class of the service.

When a RD-QoS flow with high resilience requirements (RC1) is set up, the network (additionally) reserves an alternative and disjoint explicit route for the flow with resilience requirements (see figure 1). The route reservation may be done using constraint based routing mechanisms. In case of a link or node failure, the network switches the flow to the alternative route.

For RC2, the network must reserve enough spare resources so that in the event of a failure an alternative path can be found with the required QoS. The alternate path is set up only after a failure event and its detection. To meet the required recovery time a fast failure detection in the order of milliseconds is required.

For RC3, no additional resources or alternative paths are reserved. After a failure, the network tries to recovery the affected traffic after the recovery of RC1 and RC2 is completed.

To free network resources needed for services with resilience requirements, flows of RC4 may be dropped. This will happen, if not enough spare resources are available for the recovery of RC2 to RC3 flows

Note: If no purely QoS-oriented RSVP path had been identified in advance the network has to consider the path the packets would take under non-failure circumstances before the disjoint path can be identified. This could be achieved by observing the flow of RSVP messages.

Extensions to Differentiated Services

The Differentiated Services (DS) architecture [9] realizes IP QoS by the prioritization of different services on a hop-by-hop basis. Packets are classified and conditioned at the network boundary and assigned to a behavior aggregate. The behavior aggregate is identified by bit-patterns in the DS field in the IP header, so called DS codepoints (DSCP) [10]. The DS-Field is located in the IPv4 TOS octet or IPpv6 Traffic Class octet. A specific DSCP selects a corresponding Per-Hop-Behavior (PHB) for the packet. An Expedited Forwarding (EF) PHB [11] as well as a group of Assured Forwarding (AF) PHBs [12] are already defined in RFCs with corresponding codepoints. The possible DSCPs are defined in [10].

The Network Management or a special resource control establishes a set of pre-defined routes using QoS routing in advance. In addition, the corresponding bandwidth is reserved according to the estimated or negotiated (by service level agreements) amount of traffic having resilience requirements. The packets with resilience requirements then are marked either by the application or the edge device when they enter the DiffServ network (see figure 2). In the case of a failure of either a link or a network element the network then only switches those packets to an alternative path that have the corresponding resilience bit combination set in their headers.

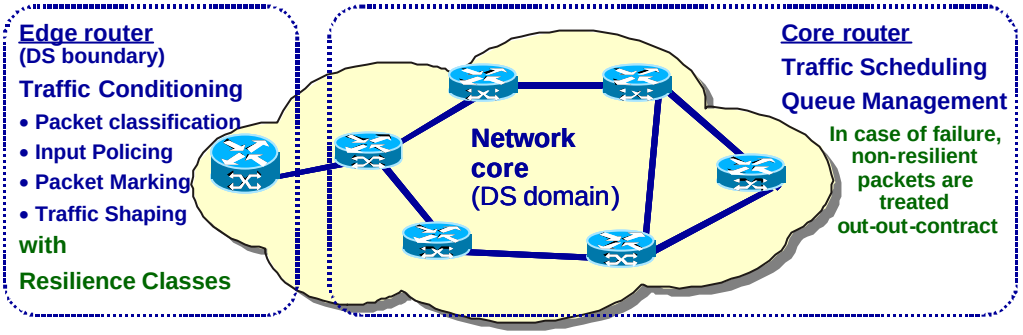

Figure 2: RD-QoS signaling with DiffServ

The marking of the packets is done using DSCP values for individual Behavior Aggregates. It is possible to define specific PHBs for behavior aggregates (BA) requiring resilience. These BAs may be independent from or extend the already defined behavior aggregates.

To allow the provisioning of end-to-end resilience over multiple administrative domains a standardized definition of the resilience behavior is needed. This may be done be using the proposed set of four resilience classes.

The bit patterns for resilience DSCPs may either be taken from the DSCP standardized pool or the pool for local and experimental use. It must however be taken care that a correct mapping of the resilience attributes over the domain borders is assured.

In case of link or node failures, the DiffServ traffic conditioner considers flows for which no or low resilience requirements were signaled (RC 3 and RC 4) to be out-of-profile. These flows are either dropped at the boundary router, or they are assigned a low drop precedence.

For flows with resilience requirements (RC 1 and 2) the required spare resources were reserved in advance. This maintains the same level of service quality in the presence network failures. Depending on the required level of resilience, the packets may be remarked and the traffic profile modified. This influences the way the traffic is shaped and when packets are dropped.

The packet classification and marking is done at the network border. Within the network packets are forwarded using standard hop-by-hop routing. The only possibility to support the 1+1 or 1:1 resilience required by RC1 is to use additionally a constraint based routing extension.

4 Integration of RD-QoS with MPLS resilience provisioning

With Multi-Protocol Label Switching (MPLS) [13] a Label Switched Path (LSP) is established between edge routers using MPLS signaling protocols. Since MPLS is path-oriented and allows a explicit route definition, various recovery mechanisms are possible [5]. The interworking of MPLS with Differentiated Services (DiffServ) allows the assignment of different

resilience levels to individual DiffServ flows [14]. An open issue is however, how to identify which DiffServ flows have to be protected and which not.

The Resilience-Differentiated QoS architecture provides a solution to this problem. The extended quality-of-service definition allows the direct mapping of RD-QoS classes to MPLS LSPs with different protection levels and recovery options according to the negotiated resilience requirements. The BA of services with high resilience requirements (RC 1) can be assigned to a LSP with a predefined dedicated protection path (1+1 protection). For service classes with medium resilience requirements, a LSP with a link or path rerouting as recovery model may be provided. For lower resilience classes, no MPLS recovery is configured. LSPs or services with no resilience requirements can be transported as low-priority, pre-emptible traffic using the protection resources of higher resilience classes when no failures are present in the MPLS domain.

Thus, the proposed architecture allows a mapping of the resilience signaling at the network edge to the provisioning of resilience mechanisms at the MPLS core network. This integrates the resilience provisioning with the QoS provisioning, so that Service Providers can flexibly offer and bill these value-added services.

5 Conclusions

New internet services with increased resilience requirements are already being offered or currently emerging. Real-time services cannot tolerate long outages, and global e-commerce and mission critical internet services require a maximum of availability and a minimum of network outage times.

Traffic Engineering methods which allow the provisioning of network resilience are a clear requirement for the future Internet architecture. MPLS is an example, where such requirements are already taken into account for the development of a new forwarding protocol, and several recovery mechanisms using MPLS are already proposed. Ongoing work in the IETF investigates the integration of MPLS with DiffServ and IntServ. With MPLS supporting DiffServ for example, it is possible to assign different levels of protection to individual flows. However, an ISP or NS is faced with the problem, which flows to protect against network failures.

This document proposes the extension of the Quality of Service signaling to include resilience requirements of IP services, thus allowing a differentiated resilience for individual services. The resilience of a service can be tailored to the actual requirements of the individual applications. This results in a more effective resource usage. Moreover, the resilience provisioning can be managed using a single administrative platform, thus reducing management complexity and operational cost. The immediate advantage for an ISP is, that the resilience can be treated as an value-adding service, which can be charged for. Finally, the proposed RD-QoS architecture with defined resilience classes allows the signaling of end-to-end resilience and QoS for IP services over multiple domains employing IntServ/RSVP, DiffServ and MPLS.

6 References

[1] D. Awduche, " MPLS and traffic engineering in IP networks ", IEEE Communications Magazine, Volume: 37, No.12, Page(s): 42 – 47, December 1999.

[2] K. Owens, M. Oommen, "Network Survivability Considerations for Traffic Engineered IP Networks", Work in Progress, Internet Draft, <draft-owens-te-network-survivability-00.txt>, March 2000.

[3] A. Autenrieth, A. Kirstädter: "Provisioning of Differentiated IP Resilience and QoS - An Integrated Approach". ITG Workshop "IP in Telekommunikationsnetzen", Bremen, Germany, January 25 - 26, 2001

[4] A. Autenrieth, A. Kirstädter: "Fault-Tolerance and Resilience Issues in IP-Based Networks". Second International Workshop on the Design of Reliable Communication Networks (DRCN2000), Munich, Germany, April 9 - 12, 2000

[5] V. Sharma, B.M. Crane, S. Makam, K. Owens, C. Huang, F. Hellstrand, J. Weil, L. Andersson, B. Jamoussi, B. Cain, S. Civanlar, A. Chiu, "Framework for MPLS-Based Recovery", Work in Progress, Internet Draft, <draft-ietf-mpls-recovery-frmwrk-02.txt>, March 2001.

[6] R. Braden, D. Clark and S. Shenker, "Integrated Services in the Internet Architecture: an Overview", RFC 1633, June 1994.

[7] R. Braden (Ed.), L. Zhang, S. Berson, S. Herzog, S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, September 1997.

[8] J. Wroclawski, "The Use of RSVP with Integrated Services", RFC 2210, September 1997.

[9] D. Black, S. Blake, M. Carlson, E. Davies, Z. Wang, W. Weiss, "An Architecture for Differentiated Services", RFC 2475, December 1998.

[10] K. Nichols, S. Blake, F. Baker, D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, December 1998.

[11] V. Jacobson, K. Nichols, and K. Poduri, "An Expedited Forwarding PHB", RFC 2598, June 1999.

[12] Heinanen, J., Baker, F., Weiss, W. and J. Wroclawski, "Assured Forwarding PHB Group", RFC 2597, June 1999.

[13] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, January 2001.

[14] F. Le Faucheur, L. Wu, B. Davie, S. Davari, P. Vaananen, R. Krishnan, P. Cheval, J. Heinanen, "MPLS Support of Differentiated Services", Work in Progress, Internet Draft, <draft-ietf-mpls-diff-ext-08.txt>, February 2001.