

ITU-T Recommendation Y.2174 (Y.RACF-DistribMPLS)

Distributed RACF architecture for MPLS networks

All rights reserved. No part of this publication may be reproduced, by any means whatsoever, without the prior written permission of ITU

Contents

1. Scope.....	3
2. References	3
3. Definitions	4
4. Abbreviations and acronyms.....	4
5. Conventions	5
6. Distributed RACF architecture and supporting requirements.....	5
6.1 – Distributed TRC-FE architecture in MPLS-TE core networks.....	6
6.2 – Requirements to support distributed architecture	9
7. Admission control process with distributed architecture.....	9
8. Security considerations.....	13

ITU-T Recommendation Y.2174 (Y.RACF-DistribMPLS)

Distributed RACF architecture for MPLS networks

Summary

The architectural structure of the RACF transport resource control functional entity (TRC-FE) is considered in this Recommendation from a distributed perspective whereby MPLS transport resource information is available in the label edge routers (LER) and gateways connected to the LERs. This Recommendation defines the distributed architecture for the TRC-FE, specifies supporting requirements, and describes the resource reservation process for this architecture.

1. Scope

ITU-T Recommendation Y.2111 [Y.2111] defines general requirements for resource and admission control functions for Next Generation Networks (NGN). These requirements are independent of the underlying transport technology.

Multi-protocol label switching (MPLS) is considered to be a key transport technology in core networks. In particular, MPLS with traffic engineering (MPLS-TE) capabilities provide significant assurance in the delivery of desired quality of service (QoS) for a variety of services and applications.

Resource and admission control functions (RACF) control the flow of traffic from access networks into the MPLS core [Y.2111]. The architectural structure of one RACF functional entity, the transport resource control functional entity (TRC-FE), is considered in this Recommendation from a distributed perspective whereby MPLS transport resource information is available in label edge routers (LER) and gateways connected to the LERs. This Recommendation defines the distributed architecture for the TRC-FE and specifies supporting requirements.

2. References

The following ITU-T Recommendations and other references contain provisions which, through reference in this text, constitute provisions of this Recommendation. At the time of publication, the editions indicated were valid. All Recommendations and other references are subject to revision; users of this Recommendation are therefore encouraged to investigate the possibility of applying the most recent edition of the Recommendations and other references listed below. A list of the currently valid ITU-T Recommendations is regularly published. The reference to a document within this Recommendation does not give it, as a stand-alone document, the status of a Recommendation.

[Y.2012] – “Functional requirements and architecture of the NGN”, ITU-T Recommendation Y.2012, September 2006.

[Y.2111] – “Resource and admission control functions in Next Generation Networks”, ITU-T Recommendation Y.2111, September 2006.

[Y.2171] – “Admission control priority levels in Next Generation Networks”, ITU-T Recommendation Y.2171, September 2006.

[Y.2172] – “Service restoration priority levels in Next Generation Networks”, ITU-T Recommendation Y.2172, June 2007.

[RFC3175] – “Aggregation of RSVP for IPv4 and IPv6 Reservations”, IETF RFC 3175, September 2001.

[RFC3209] – “Extensions for RSVP for LSP Tunnels”, IETF RFC 3209, December 2001.

[RFC3564] – “Requirements for Support of Differentiated Services-aware MPLS Traffic Engineering”, IETF RFC 3654, July 2003.

[RFC3814] – “Multiprotocol Label Switching (MPLS) Forwarding Equivalence Class To Next Hop Label forwarding Entry (FEC-To-NHLFE) Management Information Base (MIB)”, IETF RFC 3814, June 2004.

[RFC4090] – “Fast Reroute Extensions to RSVP-TE for LSP Tunnels”, IETF RFC 4090, May 2005.

[RFC4125] – “Maximum Allocation Bandwidth Constraint Model for DiffServ-aware MPLS Traffic Engineering”, IETF RFC 4125, June 2005.

[RFC4126] – “Maximum Allocation with Reservation Bandwidth Constraint Model for DiffServ-aware MPLS Traffic Engineering”, IETF RFC 4126, June 2005.

[RFC4127] – “Russian Dolls Bandwidth Constraint Model for DiffServ-aware MPLS Traffic Engineering”, IETF RFC 4127, June 2005.

[RFC4804] – “Aggregation of Resource Reservation Protocol (RSVP) Reservations over MPLS-TE/DS-TE Tunnels”, IETF RFC 4804, February 2007.

3. Definitions

For the consistency between RACF, this Recommendation uses the same definitions as defined in Y.2111.

4. Abbreviations and acronyms

BG-FE: Border Gateway Functional Entity

CAC: Connection Admission Control

CPE: Customer Premises Equipment

CPN: Customer Premise Network

DSCP: DiffServ Code Point

DS-TE: Diffserv-aware MPLS Traffic Engineering

EF: Expedited Forwarding

IP: Internet Protocol

LDP: Label Distribution Protocol

LER: Label Edge Router

LSP: Label Switched Path

MAM: Maximum Allocation Model

MAR: Maximum Allocation with Reservation

MPLS: Multi-Protocol Label Switching

MPLS-TE: MPLS Traffic Engineering

NAPT/FW: Network Address and Port Translation/Firewall

NGN: Next Generation Networks

- PD-FE: Policy Decision Functional Entity
- PE-FE: Policy Enforcement Functional Entity
- QoS: Quality of Service
- RDM: Russian Dolls Model
- RIR: Resource Initiation Request
- RIP: Resource Initiation Response
- RSVP: Resource Reservation Protocol
- RSVP-TE: RSVP Traffic Engineering
- TE: Traffic Engineering
- TRC-FE: Transport Resource Control Functional Entity
- VoIP: Voice over IP

5. Conventions

None

6. Distributed RACF architecture and supporting requirements

A high level view of the role of RACF in MPLS core networks is depicted in Figure 1. Note that this figure is logical. It does not preclude the case where the service stratum spans across multiple service provider domains. Real-time traffic is aggregated by access networks into the MPLS core; there may be multiple core networks – also referred to as autonomous administrative domains. Inside the MPLS core, label edge routers are assumed to be connected via pre-provisioned label switched paths (LSP) or traffic engineering (TE) tunnels. The terms LSP and TE tunnels are synonymous in this document. A brief overview of MPLS-TE capabilities to support admission control is provided in Appendix 1.

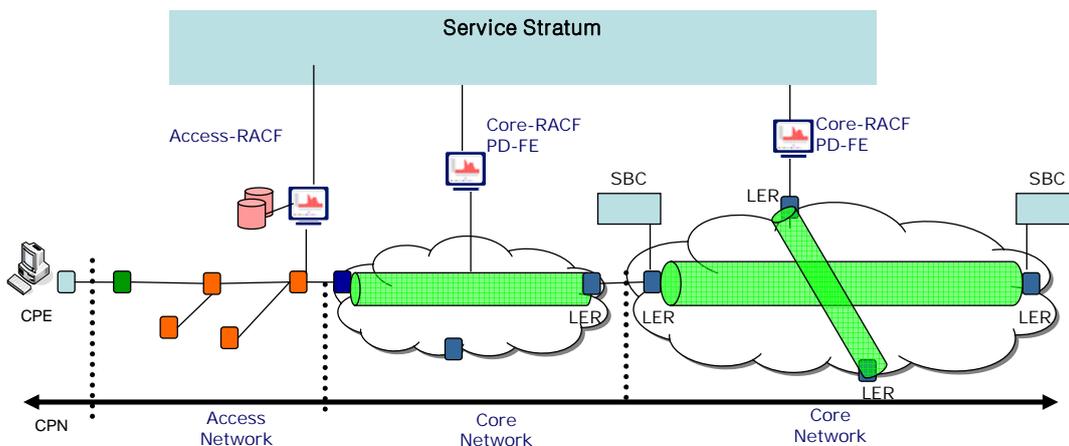


Figure 1 – Network architecture overview

Admission of traffic flows into these tunnels requires careful attention in order to achieve the required QoS. The admission control process needs to determine the type of incoming service/application flows (real-time calls, interactive sessions, or data flows) as well as their QoS and bandwidth requirements. The initial point of contact for incoming flows is the border gateway functional entity (BG-FE) [Y.2012] which is connected to the LER. The BG-FE is defined as a key node for support of dynamic QoS, network address and port translation/firewall (NAPT/FW) control and NAPT traversal. The BG-FE resides at the boundary between an access network and a core network or between two core networks. This Recommendation recommends the use of a specific implementation of the BG-FE, namely the session border controller (SBC) [b-Y.Sup2]. For the purpose of this document, the SBC is thus synonymous with the BG-FE – all SBC requirements defined here for admission control purposes are also applicable to generic BG-FE's.

The processing of incoming flows commences at the SBC prior to admission control processing at the policy decision functional entity (PD-FE). The PD-FE needs to interact with the service/application layer to determine the following QoS requirements about incoming calls/sessions:

1. Terminating SBC
2. Priority requirements for connection admission control (CAC) [Y.2171] and restoration [Y.2172], [RFC4090]
3. Bandwidth requirement (based on appropriate codec negotiation between originating and terminating SBC as well as call/session QoS requirements if any)
4. DiffServ code point (DSCP) based on the type of call/session and priority requirements.

The PD-FE then needs to relay these requirements to the TRC-FE which then proceeds to determine the availability of required bandwidth.

From an architectural perspective, the PD-FE could be developed into a single centralized entity or alternately, some PD-FE functions could be distributed in other physical elements such as the control plane in the SBC. The precise nature of the PD-FE architecture and the resulting reference point signalling flows with the TRC-FE and PE-FE are for further study. In this Recommendation, the PD-FE is depicted as a single functional entity (see Figure 2) regardless of the ultimate form of the PD-FE architecture.

Once the QoS and priority requirements have been established by the PD-FE, the flows are directed to the LER where they are mapped into the proper tunnels. Information on tunnel status and bandwidth utilization conditions is assumed to be available at the LERs and the connected gateways – tunnel selection for incoming flows is based on this information.

The architectural structure of the TRC-FE is considered in this Recommendation from a distributed perspective whereby MPLS transport resource information is available in LER and gateways connected to the LERs. This clause defines the distributed architecture for the TRC-FE and specifies supporting requirements.

6.1 – Distributed TRC-FE architecture in MPLS-TE core networks

A distributed implementation for the TRC-FE is shown in Figure 2. It is based on the capabilities of the MPLS LER and the SBC which serves as the gateway connected to the LER.

IETF RFC 3175 [RFC3175] defines a process by which gateways connected to LERs in an MPLS-TE network, are involved in the admission control of real-time applications into pre-established MPLS-TE tunnels between the LER pairs. IETF RFC 4804 [RFC4804] extends this

to the case where the TE tunnels support dynamic bandwidth allocation via DiffServ-aware MPLS traffic engineering (DS-TE). In an NGN environment, a session border controller provides key functionality as a “gateway” for real-time applications [b-Y.Sup2] and these SBC’s are connected to the LERs. SBC functions are defined as signaling path functions (e.g., signaling protocol translation and inter-working, authentication, authorization, and accounting, session-based routing) and media path functions (e.g., policing and marking, resource and admission control, opening and closing of pinhole/firewall). Media path functions such as policing and marking and opening and closing of pinhole/firewall provide policy enforcement functionality for incoming flows. The SBC therefore serves the role of policy enforcement functional entity (PE-FE) in this architecture.

The media path resource and admission control or CAC function also provides an initial level of TRC-FE functionality. The CAC function and LER are required to be RSVP-enabled while the LER is required to be RSVP-TE enabled. The CAC function is also required to keep track of available reserved bandwidth with other SBCs in the MPLS-TE core. Thus decision making for additional bandwidth reservations is carried by this CAC function based on information regarding current bandwidth states between pairs of SBCs. The CAC function and associated requirements defined below is described here as part of the SBC media path. However, it should be noted that such a CAC function can be implemented in any generic gateway or border element, rather than being limited as an SBC media path function.

Tunnel states are maintained by the LERs in the core. The LER is part of the autonomous system domain and it receives the entire route and link status information via the IP and MPLS routing protocols. LERs are required to be RSVP enabled in order to establish TE tunnels and reserve bandwidth in these tunnels [RFC3175]. Thus, they maintain the state of all originating and terminating TE tunnels. State information includes tunnel attributes as well as tunnel bandwidth characteristics (available bandwidth, bandwidth upper bound, etc). In this role, the LER provides the next layer of TRC-FE functionality.

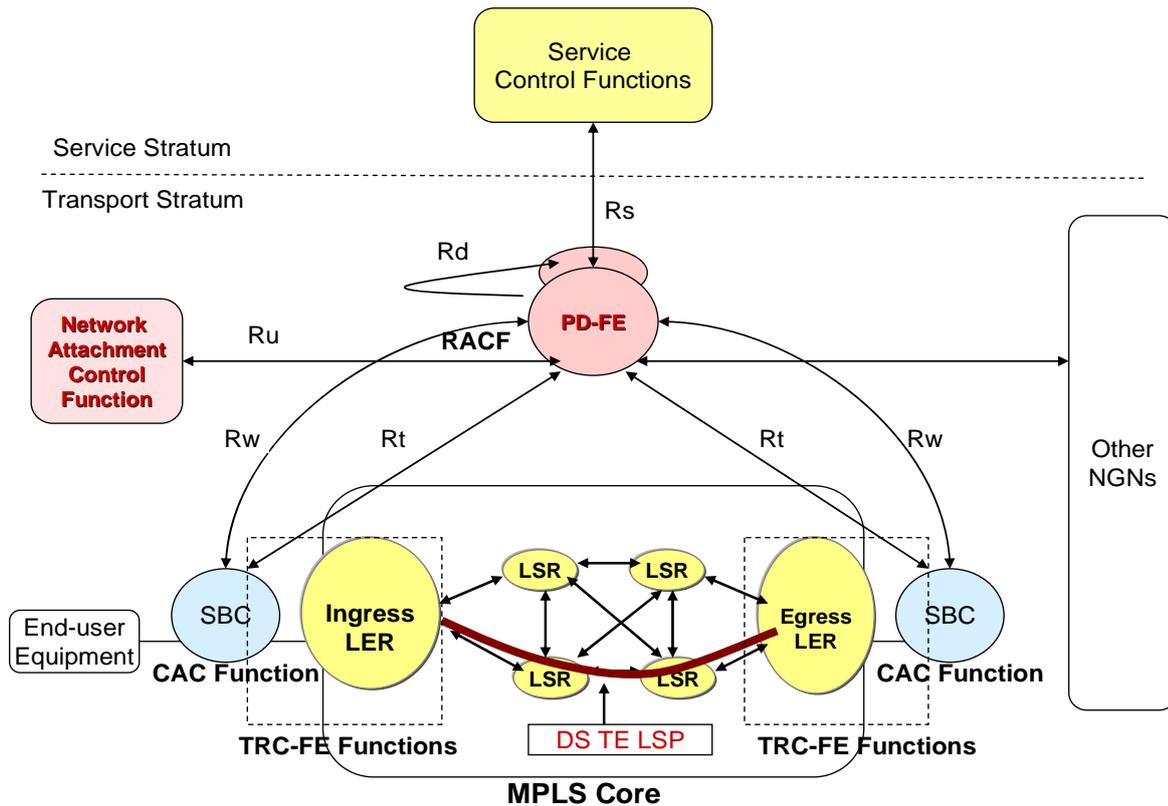


Figure 2 - Distributed TRC-FE architecture

The R_c interface functionality [Y.2111] is already provided at the LER by virtue of supporting IP and MPLS protocols as well as TE tunnel state information. It is also supported by requiring the SBC CAC functions to keep track of bandwidth status between all SBC pairs. Further, the LERs also play the role of transport resource enforcement functional entity (TRE-FE) by virtue of their capability to establish TE tunnels and reserve bandwidth. Hence, the R_c and R_n reference points [Y.2111] are not required in this implementation. However, the R_t interface is required between the PD-FE and the SBC CAC function for uni-directional and symmetric bi-directional flows.

Since this MPLS-specific transport knowledge is distributed across all LER-SBC combinations in the core, this architecture is considered to be distributed. Note that it is possible for multiple SBCs to be connected to a single LER and therefore share a tunnel between LER pairs. The TRC-FE function is jointly carried out by the SBC CAC function (determination of bandwidth state with other SBCs) and the LERs (maintenance of TE tunnels between LER pairs). Once available bandwidth has been established, the SBC CAC function admits the media flow request.

Bandwidth and QoS requirements for incoming flows are submitted to the initiating SBC CAC function by the PD-FE. The PD-FE also determines the destination SBC for the call/session flow. The detailed PD-FE mechanisms for determining the destination SBC are for further study. The SBC CAC function determines if sufficient bandwidth is available between itself and the terminating SBC (see Clause 7). If so, the SBC CAC function admits the media flow. If not, then the initiating and terminating SBC CAC functions submit RSVP reservation requests for additional bandwidth prior to admitting the flow request. Bandwidth reservation can be on a per-call or aggregated basis. The respective LERs identify the specific bi-directional tunnel – based on inputs received from the PD-FE

- and then attempt to reserve additional bandwidth. If the reservation is successful, the SBC admits the media flow. A detailed description of the admission control process is provided in clause 7.

6.2 – Requirements to support distributed architecture

The SBC CAC function provides an initial level of TRC-FE functionality (tracking reserved bandwidth availability between pairs of SBCs). The following requirements are necessary to support the CAC function:

1. The CAC function is required to track bandwidth reservation status between the SBC and all other SBCs in the core network. For each SBC pair, the CAC function needs to keep track of the total amount of bandwidth reserved and the current value of reserved bandwidth utilization as described in clause 7.
2. The CAC function is required to interact with the PD-FE via the Rt reference point in order to receive necessary call/session flow input such as source, destination, and DSCP. protocol requirements to support this reference point are for further study.
3. The CAC function is required to be RSVP-enabled per IETF RFC 3209 [RFC3209]. This allows the CAC function in the originating SBC to initiate RSVP reservations [RFC3175], [RFC4804] to the destination SBC CAC function.

The LER provides an additional level of TRC-FE function in this architecture by establishing TE tunnels and maintaining tunnel status information which is necessary to determine further RSVP reservations. The following requirements are necessary to support the LER:

1. The LER is required to be MPLS-TE enabled per IETF RFC 3564 [RFC3564]. If dynamic bandwidth allocation is desired, then the LER is required to be DS-TE enabled. In this case, all LERS in the core are required to be enabled with an appropriate bandwidth constraint algorithm. Three choices are available: maximum allocation model (MAM) [RFC4125], maximum allocation with reservation model (MAR) [RFC4126], and russian dolls model (RDM) [RFC4127]. Note that all LERs in the administrative domain are required to have the same bandwidth constraint algorithm.
2. The LER is required to map the flow into the proper tunnel according to the process described in IETF RFC 3814 [RFC3814].
3. The LER is required to be RSVP-enabled such that bandwidth reservations in the MPLS-TE/DS-TE tunnels are supported as described in IETF RFC 3175 [RFC3175] and RFC 4804 [RFC4804] as appropriate. Note that bandwidth reservations can be on a per-call/session basis or on an aggregated basis – the choice is up to the network operator.

7. Admission control process with distributed architecture

It is assumed that MPLS-TE or DS-TE tunnels are pre-established between pairs of LERs in the MPLS-TE core network and that some pre-determined amount of bandwidth is reserved in these tunnels – this is appropriate in the case where aggregated RSVP reservations is the practice. Note that additional bandwidth reservation in these tunnels may be permitted up to a specified limit.

Incoming calls/sessions are signaled to the PD-FE via the SBC. The following information and requirements about these calls/sessions are determined via appropriate interactions between the PD-FE and the Service/Application Layer:

1. Terminating SBC
2. CAC priority
3. Bandwidth requirement (based on appropriate codec negotiation between originating and

terminating SBC as well as call/session QoS requirements if any)

- DiffServ code point (DSCP) based on the type of call/session and its priority requirements.

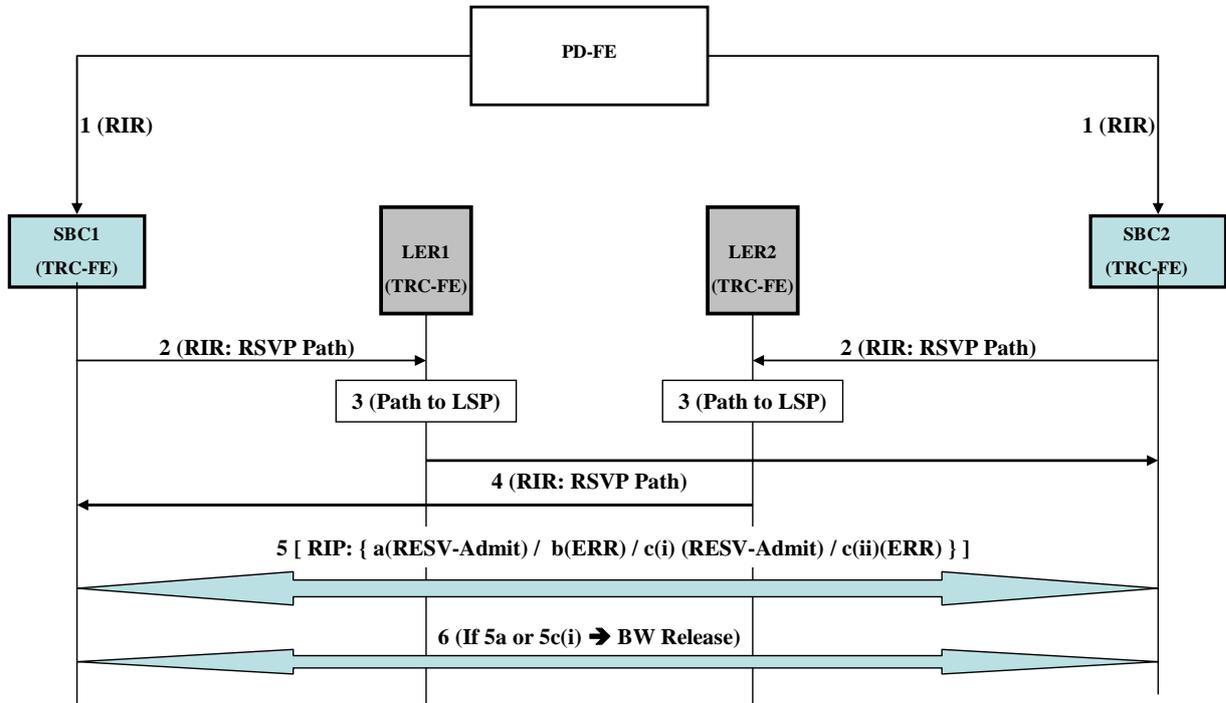


Figure 3 – TRC-FE process for per-call reservations

The next step is the determination of available bandwidth between the originating and terminating SBCs to support the flow. This determination can lead to per-call bandwidth reservation or aggregated bandwidth reservation [RFC3175].

Per-call reservation provides optimal use of bandwidth. However, if the volume of busy-hour calls/sessions is significant, then the resulting RSVP processing may create an unacceptable burden on the LERs. Hence, per-call reservations should be considered only in cases where busy-hour call/session volume is not excessive.

Assuming that no bandwidth is reserved in the TE tunnels during pre-establishment, the per-call/session bandwidth reservation is as follows (Figure 3):

1. PD-FE submits call/session input as described above to the CAC functions in the originating and terminating SBCs. Per RACF terminology [Y.2111], this is the resource initiation request (RIR). Head-end and tail-end LERs identify appropriate TE tunnel based on this input.
2. Originating SBC CAC function submits the RIR in the form of an RSVP PATH message. A similar RSVP PATH message is submitted by the terminating SBC CAC function to the tail-end LER to support the call/session in the opposite direction.
3. Head-end LER processes RSVP PATH message and determines the amount of requested bandwidth in the identified TE tunnel. Similar action is carried out by the tail-end LER.
4. The RSVP PATH messages are forwarded in both directions by the head-end and tail-end LERs to the terminating SBC and originating SBC.

5. The decision making is as follows:

- a. If sufficient bandwidth is available in the TE tunnel (both directions) to support the call/session, then the bandwidth is reserved in both directions and resource initiation responses (RIP) [Y.2111] in the form of RSVP RESV messages are sent back to the originating and terminating SBCs and the media flow is admitted in both directions.
- b. For the case of fixed bandwidth TE tunnels, if sufficient bandwidth is not available, then an RIP in the form of an RSVP Error Code message is sent back to the initiating SBC and the call/session is denied admission.
- c. For the case of DS-TE tunnels, if sufficient bandwidth is not available, then the DS-TE bandwidth constraint algorithm (MAM, MAR, or RDM) attempts to increase the bandwidth upper limit in the tunnel by adjusting bandwidth limits from DS-TE tunnels of lower pre-emption priority [RFC3564]:
 - i. If the DS-TE tunnel bandwidth limit is successfully increased, then the desired bandwidth to support the incoming call/session is reserved in both directions, RSVP RESV messages are sent back to the originating and terminating SBCs, and the media flow is admitted in both directions.
 - ii. Otherwise, an RSVP Error Code message is sent back to the initiating SBC and the call/session is denied admission.

6. Once the call/session is completed, the reserved bandwidth is released in both directions.

When busy-hour call/session volume is expected to be high, aggregated RSVP reservations are recommended. By reserving pre-determined chunks of bandwidth that can support a large number of simultaneous flows, RSVP processing on the LERs is significantly minimized; however bandwidth utilization may not be optimal.

To support aggregated reservations, the SBC CAC function is assumed to have the necessary knowledge to make the reservation decision. The SBC CAC function is required to keep track of the total amount of reserved bandwidth R and the amount of utilized bandwidth by calls/sessions in progress U between itself and all other SBCs in the network.

The detailed process for aggregated bandwidth reservations for an incoming call/session is as follows (Figure 4):

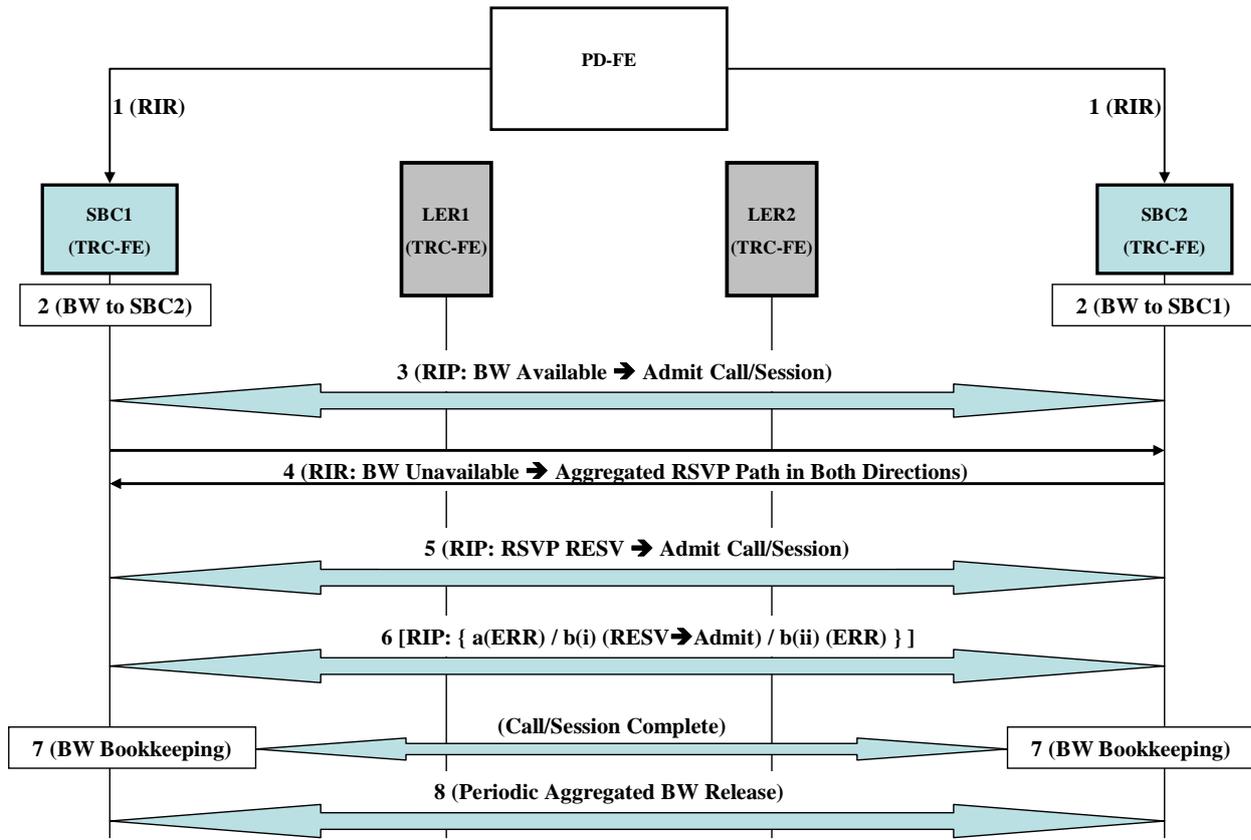


Figure 4 – TRC-FE processing for aggregated reservations

1. PD-FE submits call/session input (as described above) RIR to the CAC functions in the originating and terminating SBCs.
2. Originating and terminating SBC CAC functions compare the available reserved bandwidth $R - U$ with the bandwidth requirements of the incoming call/session.
3. If the available reserved bandwidth is sufficient (e.g., $(R-U)/R > x$, where x is some percent increase threshold), then the originating and terminating SBCs admit the media flow in both directions. The LERs map the flow into the appropriate TE tunnels based on the signaled information. The originating and terminating SBCs update the respective U values for current bandwidth utilization between them.
4. If the available reserved bandwidth is not sufficient, then the originating and terminating SBC CAC functions request an aggregated chunk of bandwidth reservation in the TE tunnels via RIRs in the form of RSVP PATH messages.
5. If bandwidth reservation is successful, RIP responses in the form of RSVP RESV messages are generated and the media flow is admitted in both directions. The two SBCs update their R and U values.
6. If bandwidth reservation is not successful, then:
 - a. For the case of fixed bandwidth MPLS-TE tunnels, an RIP in the form of an RSVP Error Code message is sent back to the initiating SBC and the call/session is denied admission.

- b. For the case of DS-TE tunnels, the DS-TE bandwidth constraint algorithm (MAM, MAR, or RDM) attempts to increase the bandwidth upper limit in the tunnel by adjusting bandwidth limits from DS-TE tunnels of lower pre-emption priority [RFC3564]:
 - i. If the adjustment is successful, the RSVP aggregated reservation process commences (RSVP RESV messages submitted), and the media flows are admitted in both directions. The R and U values are then adjusted accordingly between the two SBCs.
 - ii. If the adjustment is not successful, an RSVP Error Code message is sent back to the two SBCs and the call/session is denied admission.
7. Upon completion of the call/session (Case 3, Case 5, or Case 6b(i)), the respective SBC CAC functions need to adjust their reserved bandwidth utilization values (U) accordingly.
8. Periodically, if the reserved bandwidth is determined to be excessive with a significant amount of unutilized bandwidth, per pre-determined rules, chunks of bandwidth can be released during low utilization periods (e.g., $(R-U)/R > y$, where y is some percent decrease threshold). The respective SBCs need to adjust their R values accordingly. The threshold for releasing reserved bandwidth y may be different from that for increasing bandwidth reservation x .

8. Security considerations

General RACF security considerations are described in ITU-T Recommendation Y.2111 [Y.2111]. MPLS-TE security considerations are described in IETF RFC 3175 [RFC3175] and IETF RFC 4804 [RFC4804].

Appendix I - MPLS traffic engineering capabilities to support CAC

(This appendix does not form an integral part of this Recommendation)

QoS in the core network is enhanced by the traffic engineering capabilities of MPLS-TE. MPLS has been designed to overcome IP's least cost routing limitation and provides ATM-like QoS capabilities via traffic engineering methods. It relies on the IP routing protocols to build the network map and assigns labels to links that meet certain quality criteria via the label distribution protocol (LDP) or signaling protocols such as RSVP-TE [b-RFC2702]. Packet forwarding is based on swapping labels. Differentiated services (DiffServ) provides a QoS treatment to traffic aggregates. It is scalable and does not require per-flow signaling. However, it cannot guarantee QoS and it does not influence the path of a packet. MPLS can force packets into specific paths and, in combinations with constraint-based routing, can guarantee bandwidth for forwarding equivalency classes. But MPLS by itself cannot specify class-based differentiated treatment of packet flows.

Combining DiffServ and MPLS-based TE can lead to true QoS in IP packet backbones. To achieve this functionality, networks have to be carefully engineered with traffic engineering applied on a per-class basis – DiffServ-Aware MPLS traffic engineering (DS-TE) [RFC3564]. The goal of DS-TE is to guarantee bandwidth separately for critical traffic types (e.g., emergency telecommunications) such that the QoS requirements compliance for that traffic type is improved and optimized. It is assumed that the majority of traffic is of the “best effort” traffic class. Under congestion or failure conditions, this traffic class will have reduced bandwidth available in order to ensure that more critical traffic classes have the required bandwidth to meet their priority and QoS requirements. An aggregated grouping of traffic trunks based on the class of service requirements such that they share the same bandwidth reservation is called class type (CT). Up to eight class types are allowed. Each CT has two attributes:

- Bandwidth constraint (BC) – A limit on the percentage of a link's bandwidth that a particular CT may request. Three bandwidth constraint models have been developed: maximum allocation model (MAM) [RFC4125], maximum allocation with reservation model (MAR) [RFC4126], and russian dolls model (RDM) [RFC4127]. Note that each MPLS network domain can be provisioned to support only one of these BC models. All tunnels in a domain have to be governed by the rules of one model only.
- Preemption priority (p) – The relative importance of a given CT compared to others. This priority enables the DS-TE bandwidth constraint models to release shared bandwidth occupied by a lower priority CT when higher priority CT traffic arrives during conditions of congestion or failure.

Additional attributes associated with incoming flows are as follows:

- DiffServ code point (DSCP): The desired QoS of an incoming flow can be characterized by delay, packet loss, and jitter requirements. These requirements can be summarized by assigning appropriate DSCP values to the media stream. For example, stringent delay, loss, and jitter requirements are characteristics of voice over IP (VoIP) services. These services can then be characterized by the expedited forwarding (EF) DSCP as this code point refers to stringent QoS requirements.
- Restoration/re-route priority: This is the priority with which a tunnel can be restored [Y.2172] or mapped into an MPLS fast re-route priority [RFC4090].
- MPLS EXP Bit: The DSCPs are mapped to the three EXP bits in the MPLS header [b-RFC3270].

From the perspective of admission control, an incoming service or application seeking entry into the MPLS network needs to be “mapped” into the TE tunnel of the appropriate class type. Such a mapping can be done by linking the service CAC priority as defined in ITU-T Recommendation Y.2171 [Y.2171] with the priority attribute of the appropriate class type that supports the desired QoS requirements of the incoming service.

Admission control policy enforcement for unicast real-time applications (e.g., voice over IP calls, video services) may be accomplished for MPLS-TE networks based on functionality described in IETF RFC 3175 [RFC3175] and RFC 4804 [RFC4804]. Policies for multicast applications are for further study. The underlying premise is the setup of TE tunnels between pairs of LERs with sufficient bandwidth such that a significant number of calls/sessions for various applications can traverse these tunnels without additional processing of signaling messages on the intermediate routers along the path of the tunnel.

A network that is DiffServ-enabled and is RSVP-aware offers several mechanisms to support topologically aware admission control that can be per-flow control or on aggregated flows. The advantage of aggregated RSVP reservations is that it offers dynamic admission control without the per-flow reservations and associated RSVP signaling in the DiffServ core. With DS-TE tunnels, different types of services/applications can be directed into designated tunnels identified by a unique class type. Further, appropriate bandwidth allocation/constraints can be enforced for these different class type tunnels such that service/application QoS requirements are satisfactorily met.

The benefits of aggregating RSVP reservations over TE tunnels are as follows:

- RSVP signaling messages for bandwidth reservations are ignored by intermediate routers along the tunnel path per RFC 3175 [RFC3175]. RSVP message processing is further minimized by initiating them for aggregated reservations as opposed to per-call reservations.
- Core scalability is unaffected because the core has to simply maintain aggregated TE tunnels.
- Aggregate reservations can be network engineered with constraint based routing taking advantage of alternate paths when needed.
- Aggregated reservations over TE tunnels can be protected against failure via MPLS fast reroute [RFC4090].

Bibliography

[b-RFC2702] – “Requirements for Traffic Engineering over MPLS”, IETF RFC 2702, September 1999.

[b-RFC3270] – “Multi-Protocol Label Switching (MPLS) Support for Differentiated Services”, IETF RFC 3270, May 2002.

[b-Y.Sup2] - Supplement 2 to Y-series Recommendations – ITU-T Y.2012, Supplement on session/border control (S/BC) functions (S/BC).
