

A Quality-of-Service Resource Allocation Client for CASP

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of Section 10 of RFC2026.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as “work in progress.”

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

To view the list Internet-Draft Shadow Directories, see <http://www.ietf.org/shadow.html>.

Copyright Notice

Copyright (c) The Internet Society (2002). All Rights Reserved.

Abstract

Signaling resource reservations is one of the possible applications of the Cross-Application Signaling Protocol (CASP). This document describes a client protocol that supports per-flow resource reservation for unicast and source-specific multicast flows, in both in-band and out-of-band modes, in sender- and receiver-directed operation.

Contents

1	Introduction	2
2	Operation	3
3	Objects	3
3.1	Next (N)	3
3.2	Priority (P)	4
3.3	Version (V)	4
3.4	Partial Reservation (PR)	4
3.5	Status (S)	4
3.6	Time (T)	4
4	Messages	5
5	Resource Objects	5
5.1	Bandwidth (B)	5
5.2	PHB	5
5.3	RSVP Rspec	5
5.4	L2 Properties	5

6	Local Information	6
7	Route Change and Mobility Considerations	7
8	Security Considerations	7
9	Open Issues	8
10	Acknowledgements	8
11	Authors' Address	9

1 Introduction

CASP-QOS is a client protocol for the Cross-Application Signaling Protocol (CASP) [1]. It is one of a family of CASP client protocols and offers per-flow resource allocation and reservation.

CASP-QOS has the following properties:

Direction-neutral: The protocol supports both receiver-oriented and sender-oriented reservations. In each mode, the non-reserving side can suggest QoS parameters. For example, the data receiver can send the first CASP message to indicate the range of bandwidths and QoS parameters it is willing to tolerate, but the data sender makes the actual reservation within that range.

Bidirectional reservation: Bidirectional reservation refers to three different modes of operation. In the first, there is a single reservation message for both directions, i.e., a traffic selector that specifies traffic flowing in both directions, typically with reversed source and destination address and port numbers. Such reservation is only feasible if the route is symmetric. Its main advantage is atomicity, so that a reservation in the forward direction is made only if traffic in the backward direction can also be accommodated.

A second, looser form of bidirectional has messages from the originator and the destination cause reservations to be set up. As before, this requires symmetric routes for in-band signaling messages and AS-symmetric routes for per-AS reservation. As shown in [2], a majority of routes are not symmetric at the AS level.

Thirdly, a reservation from the originator may trigger an independent signaling session from the destination to the originator. This mode works even if the data path is asymmetric and requires no particular protocol support at the signaling layer.

CASP QOS supports all three modes.

Reservation range: To reduce the number of reservation message exchanges, the bandwidth object contains a lower and upper bandwidth range. Nodes attempt to reserve the highest amount of resources below the maximum and update the amount accordingly. Nodes with higher reservations than the path minimum are updated on the return path.

Partial reservation: CASP-QOS messages can indicate whether they are satisfied to obtain partial reservations, i.e., reservations that only succeed on some routers [3].

Advance reservation: CASP-QOS allows to define the start and end time of a resource reservation, to support advance reservations.

Query/reserve/commit mechanism: If desired, an end system can query for available resources, reserve them and commit them. Only committed resources can be used.

Limited multicast support: CASP-QOS, like CASP, supports source-specific multicast (SSM) [4]. It does not support receiver diversity, reservation styles, blockade states and NBMA next-hops. Note that some aspects of reservation styles can be supported by appropriate traffic selectors.

2 Operation

CASP-QOS defines five message types:

Query: The QUERY message determines if a particular path has sufficient resources, but does not reserve or commit any resources. It is generally sent as NOOP/AR. The destination responds with a RESPONSE message.

Response: The RESPONSE message does not change the state of a resource reservation, but simply reports on the result of an earlier reservation or commit operation. The message is routed in the same manner as the request, except that the route is reversed or SF converted to SB.

Reserve: The RESERVE message requests a particular reservation. It generates a RESPONSE indicating the degree of success or failure. It may request a COMMIT message. The RESERVE message can also contain a response to an earlier reservation, thus allowing bidirectional reservation with bifurcated paths in one message exchange.

Commit: The COMMIT message commits resources reserved previously with RESERVE if the reservation timestamp is the same as the original reservation. If not, it creates and commits a new reservation, removing the original one. A COMMIT message that uses an existing reservation SHOULD NOT fail. Each COMMIT message carries a BANDWIDTH object just in case it visits a node that has not seen a RESERVE before.

Release: The RELEASE message releases all resources for this session, regardless of the version. It generates a RESPONSE indicating success or failure with the Status object.

3 Objects

CASP-QOS messages can carry objects described below.

3.1 Next (N)

The Next object indicates the next request that the receiver should generate if the request itself was successful. For example, if a RESERVE message contains Next = COMMIT, the receiver of the RESERVE commits the resources just reserved.

TBD: Are there other combinations beyond Reserve/Commit? Should this be used to force bidirectional reservation?

3.2 Priority (P)

The Priority object indicates the priority of the resource request. Depending on local policy, high-priority requests may either be treated preferentially compared to those with lower priority when queueing for resources or may be allowed to preempt existing resource reservations with lower priority. This facility may be used for emergency telecommunications services. Local policy determines whether a particular user is authorized to exercise a priority level.

3.3 Version (V)

The version indication is used to quickly determine whether a traffic selector or a QoS object has changed, without having to do a bit-by-bit comparison. The indication only has to be unique within the session and does not have to be monotonically increasing by one for each version; it is typically a high-resolution timestamp.

3.4 Partial Reservation (PR)

The Partial Reservation object describes how many failed reservations are allowed before reservation attempts terminate. There are two up counters that tally the number of routers where admission control failed, including due to a failed admission control procedure, and the number of routers where reservation could not be performed since the node did not speak CASP or CASP-QOS. Two additional down counters are set by the originating node to the maximum number of routers that can fail or be indeterminate before the message fails.

A “retry” flag asks each node to retry until it succeeds, using an appropriate local algorithm. Retries are only performed if the reservation failed due to insufficient resources rather than an authentication failure.

3.5 Status (S)

The Status object is added by each CASP-QOS node. It describes the nodes identity (address), whether the resource was reserved and at what level, whether authentication failed, whether reservation failed or no reservation was attempted. The Status object is returned in a Response message. The Status object also indicates the reservation version that was used for the reservation attempt.

3.6 Time (T)

The Time object describes the time the resource reservation is to be effective, expressed as a start and ending time written in NTP time format. (TBD: One could express periodic reservations in the style of iCal [5], but the complexity seems unwarranted.) A start time of zero indicates an immediate start.

Note that the CASP state has to be maintained between the time the first message is sent requesting the reservation and the end of the reservation. A requestor SHOULD use a suitably long lifetime. TBD: Should there be a notification to the initiator at the beginning of the actual reservation that indicates a new, lower refresh interval?

For resources related to conferences, it is often insufficient to find out at the time of the conference that resources are unavailable, after much effort has been expended on agreeing on a common time. A reservation for the first available time slot seems an attractive service, but difficult to set up due to the need for coordination.

4 Messages

The table below indicates which objects MAY (O), MUST NOT (–) or MUST (M) appear in each message type.

Object	Abbr.	Query	Reserve	Resp.	Commit	Release
Version	V	M	M	M	M	–
Partial Reservation	PR	O	O	–	O	–
Status	S	–	–	M	O	O
Time	T	–	O	–	–	–
Bandwidth	B	M	M	–	M	–

5 Resource Objects

A set of objects are used to request resources. Additional objects are likely to be defined in the future. It is possible to include several such objects if it is allowed for the CASP node to satisfy any one of them, starting with the one listed first. The response indicates which resource was used.

5.1 Bandwidth (B)

The Bandwidth object contains two bandwidth values, an upper and a lower bound. Each node attempts to reserve the upper bound. If it obtains resources that are below the upper bound and above the lower bound, it updates the upper bound to that lower value. The return message then updates all reservation to the upper bound seen by the destination.

Others have proposed a loss rate or explicit delay indication, possibly with a violation probability. It is not clear that there are scheduling and admission control mechanisms that can usefully guarantee such behavior on a per-flow basis. Thus, this memo does not include them.

5.2 PHB

The PHB object requests that traffic matching the traffic selector is assigned a certain per-hop behavior, such as AF12 [6].

5.3 RSVP Rspec

The Rspec object describes delay and reserved bandwidth as for Integrated Services for RSVP [7].

TBD: Should it re-use the data description or the whole record format, bit-by-bit?

5.4 L2 Properties

The L2 object describes layer-two properties abstractly, such as “low delay, but do not care about packet losses” or “high reliability”. These requests then set up specific link layer behavior. This feature requires further study.

6 Local Information

Usually, QoS-related information is generated by a host and sent to a peer representing the opposite terminating point of the path. The data has significance all along the signalling path. CASP offers the possibility to restrict the scope of signalling information to a section of the path, e.g. to a specific AS or subnet within an AS. CASP-QoS information can be added and deleted anywhere in the network. This path is referred to as localized path. Local scoping has the advantage that QoS signalling information can be limited to certain sections of the path without the need for end-to-end transport. The localized path is determined by a source node, which adds specific QoS signalling information and a sink node, which deletes the data with local scope.

Authentication of source and sink node for the path segment is required to enable message integrity for the carried QoS signalling data.

The transport of local information is useful for a number of applications, some of which are enumerated below:

Authorization token: An authorization token is a CMS encapsulated (digitally signed or encrypted) collection of objects. Such a token can be included for example by a policy decision point to allow other CASP nodes along the path (within the same administrative domain) to execute policy based admission control securely without need to retrigger a $PEP < - > PDP$ communication. Furthermore linking authorization of protocols is an other example. Protecting information which is relevant and secured within a local domain only is possible.

DSCP (DiffServ codepoint): Specific codepoints may be used within an AS for definition of a certain per-hop behaviour (PHB). The codepoint may be set by an ingress router or by the end host. The DSCP in this situation is valid within the local AS and has to be mapped to a different value when entering the next AS. A new mapping may be required because codepoint values are used differently or there is no exact mapping of the PHB between two neighbouring ASs.

Aggregation information: Information about the aggregation of signalling state can be transferred between the aggregation ingress and the aggregation egress point. This includes information about granularity of aggregation and the role (aggregation or de-aggregation) a node should act for a specific flow

Reservation priority: Information about the reservation priority may be shared along CASP-QoS capable nodes to determine the priority of a resource reservation request in the packet forwarding plane.

Accounting: Accounting and charging information (as authorization tokens) can be distributed along the signaling path to enable the creation of proper accounting records.

Operator information: Any kind of operator information may be shared by CASP-QoS nodes along the signalling path.

There is an optional **Local Object** to carry local QoS information. Each object carries an identification and a type indicator. The scope field provides hints about the scope of the local data since the processing of local data may depend on the current scope value. There is no syntax definition for the object's data structure as part of the CASP-QoS protocol. Following this concept much flexibility is given for the syntax definition of local data, e.g the CASP-QoS protocol does not care about the structuring of accounting information within a certain AS. As well the determination of proper source and sink nodes for local data has to be handled by some mechanism other than CASP-QoS.

Nesting of local information is possible, so that within a path segment local data is transported and there is a further subsegment along the same path which nests further local information. The local data is associated with a scope level. Each CASP-QoS node depending on the value of the scope has to decide whether it should process specific local data or ignore it. With the example of nesting local information, the nested path can be associated with a higher value for the scope field than the original path. This way local data associated with the original path can transparently pass the nested path. Each local object carries data for a specific path or path segment avoiding the necessity to carry local data with different scope in a single object.

7 Route Change and Mobility Considerations

Like CASP, CASP-QoS provides a means to adapt to route change. During the CASP transport connection repair upon a route change, a CASP-QoS reserve message can be triggered before sending a CASP-ADD message in the CASP-ADD originator, or after receiving CASP-ADD message at the merging node. CASP-QoS release message can be triggered before sending CASP-DEL message. In route change cases, after CASP creates necessary M states in the new path, the CASP-QoS functionalities can be provided in the following way:

- Query: CASP-QoS can discover the resource availability information along the new path by Query/Response message.
- Reservation along the new path by Reserve/Commit message.
- Release resources along the old path by Release/ Response message.

Mobility is regarded as a special case of route change in CASP processing, which needs special processing in CASP-QoS. After detecting any route change, a C state uses the same identifier as before route change. Mobility case differs from the normal route change cases in that the identifier (the source or destination address) of the transport connection is changed from the old Care-of-Address (oCoA) to the new Care-of-Address (nCoA).

Figure 1 shows an example of the processing of mobility in CASP-QoS. Assume some time there is an CASP transport association between the mobile node (MN, as destination) and the correspondent node (CN, as source) along an access router oAR. After the MN detects that the source address of outgoing packets towards destination CN has been changed from oCoA to nCoA, it uses scout message(s) towards CN to find its access router nAR and the merging point CR. After CASP creates the transport association between the MN and CR, CASP-QoS can be triggered to query or reserve resources for its QoS client. If the CR finds the dead route removal flag is set, while releasing the CASP transport states in old path, it can subsequently triggers the Release operation for the reverse direction (from the CR to the oAR) of CASP-QoS client.

8 Security Considerations

CASP relies on the security mechanisms described in [1]. Securing the messaging layer in a CASP-peer to CASP-peer fashion is provided either by IPsec or by TLS. Non peer-to-peer protection of client layer objects is provided by CMS which allows resource objects and related objects defined in this document to be encapsulated and protected by CMS. Hence no separate specification within CASP is necessary to

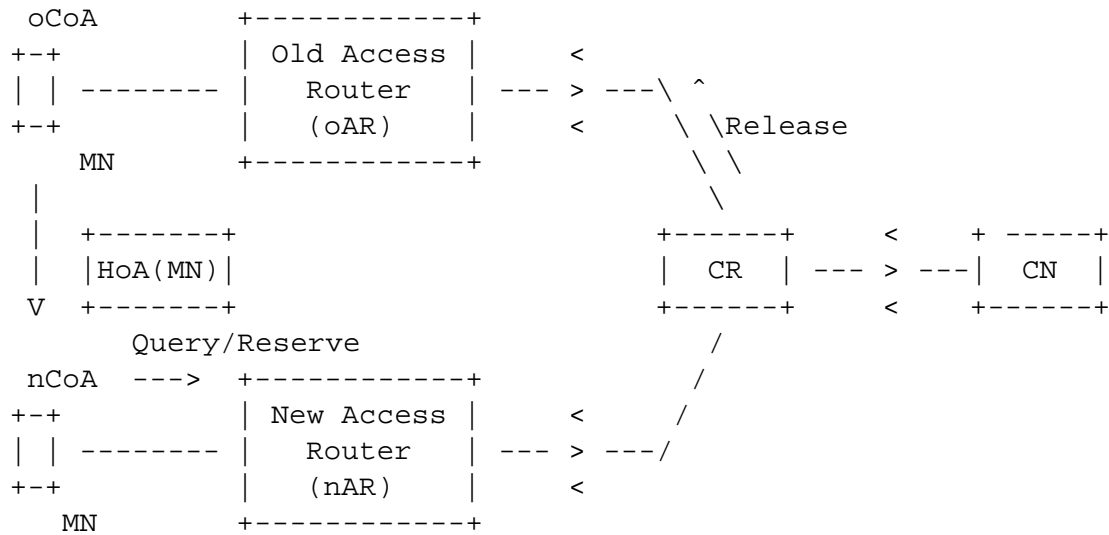


Figure 1: Mobility

describe the format of these objects. This allows some flexibility in including protected objects to link the authorization step of different protocols (for example CASP and SIP) and to transport local information within domains. The functionality described in [8] and [9] can be provided without substantial protocol modification/extensions.

9 Open Issues

- Next request feature?
- Allow multiple types of resource objects in one request?
- Usage scenarios for non-repudiation need to be described.
- Interaction with accounting/charging and related protocols (for example AAA) needs to be elaborated.
- Should there be a notification to indicate a change in refresh interval?
- The usage of authorization tokens needs to be described in more detail (message formats and an indication of useful objects).
- It might be useful to allow a traffic sink to ask the traffic source to set up a resource reservation. The sink would send a CASP request directly to the source, which would start a normal CASP reservation. For some applications, this can be done already, e.g., using SIP.

10 Acknowledgements

Robert Hancock provided useful feedback.

11 Authors' Address

Henning Schulzrinne
Dept. of Computer Science
Columbia University
1214 Amsterdam Avenue
New York, NY 10027
USA
EMail: schulzrinne@cs.columbia.edu

Hannes Tschofenig
Siemens AG
Otto-Hahn-Ring 6
81739 Munich
Germany
EMail: Hannes.Tschofenig@siemens.com

Xiaoming Fu
Technical University Berlin
Sekt. FT 5-2, Einsteinufer 25
10587 Berlin
Germany
EMail: fu@ee.tu-berlin.de

Jochen Eisl
Siemens AG
Otto-Hahn-Ring 6
81739 Munich
Germany
EMail: Jochen.Eisl@icn.siemens.de

References

- [1] H. Schulzrinne, H. Tschofenig, X. Fu, J. Eisl, and R. Hancock, "CASP - cross-application signaling protocol," Internet Draft, Internet Engineering Task Force, Sept. 2002. Work in progress.
- [2] L. Amini and H. Schulzrinne, "Observations from router-level internet traces," in *DIMACS Workshop on Internet and WWW Measurement, Mapping and Modeling*, (Piscataway, New Jersey), Feb. 2002.
- [3] P. Pan and H. Schulzrinne, "Processing overhead studies in resource reservation protocols," in *17th International Teletraffic Congress*, (Salvador da Bahia, Brazil), Sept. 2001.
- [4] H. Holbrook and B. Cain, "Source-specific multicast for IP," Internet Draft, Internet Engineering Task Force, Feb. 2002. Work in progress.
- [5] F. Dawson and D. Stenerson, "Internet calendaring and scheduling core object specification (icalendar)," RFC 2445, Internet Engineering Task Force, Nov. 1998.
- [6] J. Heinanen, F. Baker, W. Weiss, and J. Wroclawski, "Assured forwarding PHB group," RFC 2597, Internet Engineering Task Force, June 1999.

- [7] J. Wroclawski, "The use of RSVP with IETF integrated services," RFC 2210, Internet Engineering Task Force, Sept. 1997.
- [8] L. Hamer, B. Gage, and H. Shieh, "Framework for session set-up with media authorization," Internet Draft, Internet Engineering Task Force, July 2002. Work in progress.
- [9] L. Hamer, B. Gage, M. Broda, B. Kosinski, and H. Shieh, "Session authorization for RSVP," Internet Draft, Internet Engineering Task Force, July 2002. Work in progress.