Internet Engineering Task Force Working Group
Internet Draft
H. Schulzrinne
Columbia U./Siemens/U. Goettingen/Siemens
draft-schulzrinne-nsis-casp-qos-01.txt
3 March 2003
Expires: September 2003

A Quality-of-Service Resource Allocation Client for CASP

STATUS OF THIS MEMO

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 in both sender- and receiver-directed modes operation.

1 Introduction

CASP-QOS is a client protocol for the Cross-Application Signaling Protocol (CASP) [1]. It is one of a

family of CASP signaling client protocols (NSLPs) 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 accomodated.

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 initiator may trigger an independent signaling session from the responder to the initiator. This mode works even if the data path is asymmetric and requires no particular protocol support at the CASP messaging 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].

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.

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.

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.

Success: The SUCCESS response message indicates the requested reservation has been succeeded (not needed if the destination returns a Commit in response to a Reserve); includes sequence number of Reserve or Commit.

Error: The ERROR response message indicates reservation has been failed; may also release all resources already made for this session. It includes a sequence number of the failed Reserve or Commit message, and the node address where the Reserve or Commit request failed.

There are two types of (QoS NSLP) states: Reserve state and Commit state, established and refreshed by Reserve and Commit messages, respectively. An Error message may remove both states, if a "removal" indicator is provided.

Two examples for CASP-QoS operations are shown in Figure 1, where (a) shows an example for sender-directed operations, while (b) illustrates a possible handling of reservation failure and partial reservations.

3 Objects

CASP-QOS messages can carry objects described below.

3.1 Version (V)

The version indication is used to quickly determine which Resource object is used in the session, and what semantics should be assumed for that Resource object (FlowSpec).

3.2 Partial Reservation (P)

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.

4 Messages

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

Definitions of CASP-QoS message formats are give in Appendix A.

Object	Abbr.	Query	Reserve	Success	Commit	
Release	Error					
Version	V	M	M	M	--	M
Partial Reservation	P	O	O	--	O	-- M
Bandwidth	B	M	M	--	M	-- O

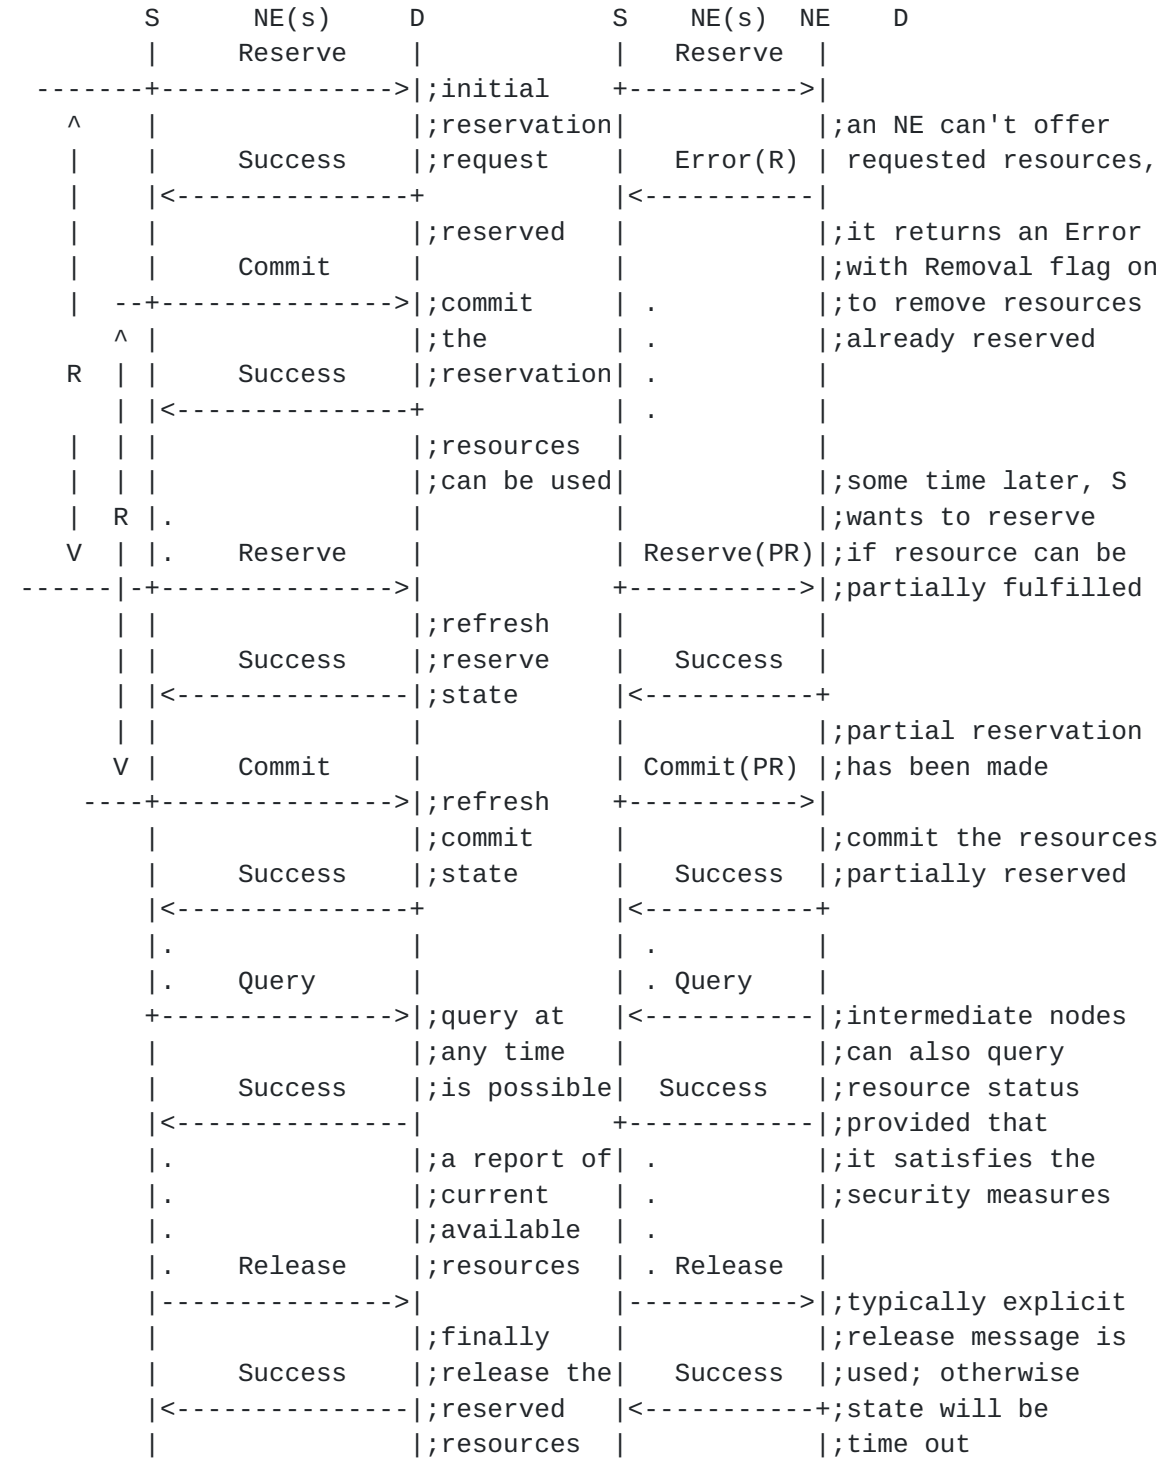

Figure 1: Two CASP-QoS Examples

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 [4].

5.3 IntServ Flowspec

The IntServ Flowspec object describes reserved resources for Integrated Services [5].

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

The separation between M-layer state and C-layer state in CASP is logically (and not physically). Hence a route change also requires to create a new reservation along the new path until the merge point (cross-over router) is reached. The separation between the Session ID and the Traffic Selector enables the merge point to associate an existing reservation with the Session ID provided by the incoming signaling message. As a difference between a standard route change and a mobility scenario the possibility of a Traffic Selector change should be mentioned. Typically the former does not require signaling messages to be forwarded beyond the merge point. The situation might be different in case of mobility and the used Traffic Selector. Thus there is an interaction with a micro/macro-mobility scheme. Using a micro-mobility scheme which is able to trigger CASP would therefore further limit double reservations and would speed the reservation setup time.

A signaling message initiator can request the deletion of the old reservation (along the old path). deleting a reservation along an old path might not always be desired. The initiator is thereby a CASP node which detects the route change first. There are some reasons why such a behavior might not be necessary or desired (for example bi-casting of

data traffic to both access routers).

Particularly of interest for Candidate Access Router (CAR) discovery is the QUERY message which allows to discover the resource availability information (and possibly reservations costs) along new candidate paths.

Subsequently a few basic mobility signaling message exchanges are described. The first exchange covers a QoS reservation for upstream traffic. Subsequently the implication for downstream traffic is explained. Thereby the interworking with micro-mobility schemes is briefly described based on Context Transfer and Hierarchical Mobile IP.

Figure 2 shows a message flow for an upstream reservation in CASP-QOS. Initially a mobile node establishes state information (and a QoS reservation) between the old access router (oAR) and the CN via several routers along the path. The mobile node has obtained a care-of address (oCoA) and might additionally have additional IP addresses for mobility (e.g. home address). The implications of selecting a Traffic Selector are discussed at the end of this section.

As soon as the mobile node detects a route change due to mobility (e.g. based on a layer 2 trigger) mobility related protocols are triggered. A CASP signaling message is required either triggered by the refresh timer or by a mobility management component at the mobile node. The signaling message establishes new state and a new reservation at the new access router (because no existing state is found for the provided Session Identifier). The signaling message is then forwarded until it reaches the cross-over router (CR). The cross-over router is identified automatically since it is the first node along the path where state information identified with the provided Session ID exists. Security issues (referred as Session/Reservation Ownership) are covered in [1] since they are independent of a particular client layer protocol.

As previously mentioned it is possible (and sometimes desired) to remove the old reservation. A dead-branch-removal flag allows the cross-over router to trigger a RELEASE message along the old path. The signaling message is then forwarded towards the CN. How far to forward depends on the used Traffic Selector and on the micro-/macro-mobility scheme.

If reservations along the old path are not released then they time out (soft-state principle). The lifetime of a reservation depends on the selected refresh interval (lifetime) which is allowed to vary between peers in the CASP chain.

Figure 3 shows a downstream reservation. Without mobility protocol interaction data packets would simply follow the new path toward the new care-of address. Hence the protocol interaction can be considered as a

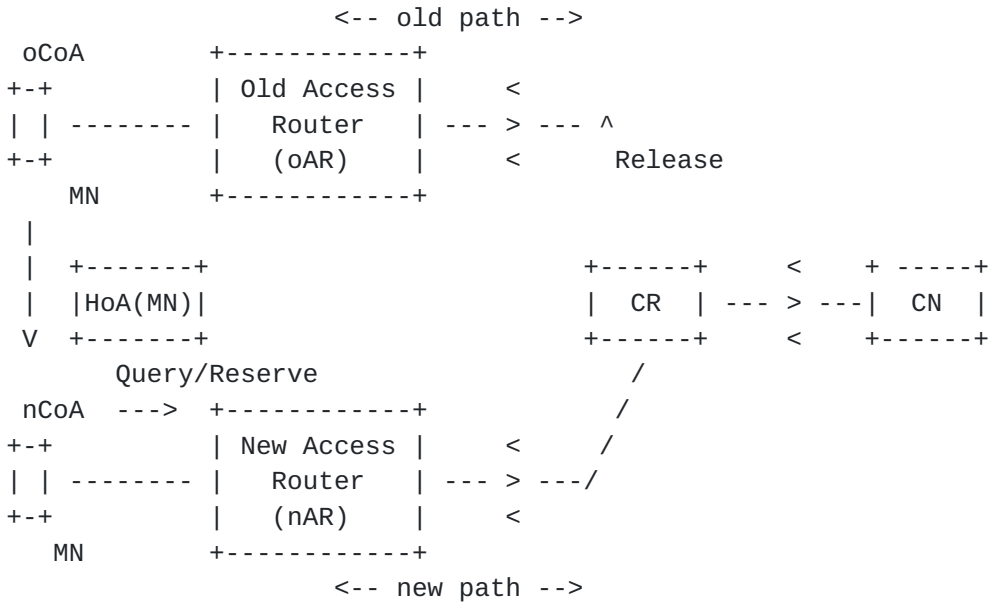

Figure 2: Upstream Reservation due to Mobility

regular route change behavior. An interaction with a mobility protocol would however enable a performance improvement. When a mobile node transmits a Binding Update/Registration Request to update its mobility binding for example at the MAP (in case of Hierarchical Mobile IP) then such a message could be used to trigger a downstream reservation. As described in the previous example the old reservation can also be removed if desired. It is worth noting that the signaling message exchange described in 2 cannot be used to trigger a downstream reservation immediately since the two paths (and the cross-over point for the upstream and the downstream traffic) might be different.

Figure 4 shows the implications for CASP when Context Transfer (CT) is used. CT allows QoS and security state established at the old access router to be forwarded to the new access router. Although various performance improving techniques would be possible, a generic approach would be to trigger the Context Transfer procedure and then to send a CASP CREATE message to the new access router. Note that the new access router could transmit this message on behalf of the mobile node. Depending on the micro-mobility scheme and on the routes used by the data traffic either a cross-over router appears somewhere along the path (e.g. router R1) or in case of host-specific routes (and tunnels) which forward traffic from the old to the new access router. CT would also provide a simplification to the security aspects in the following three

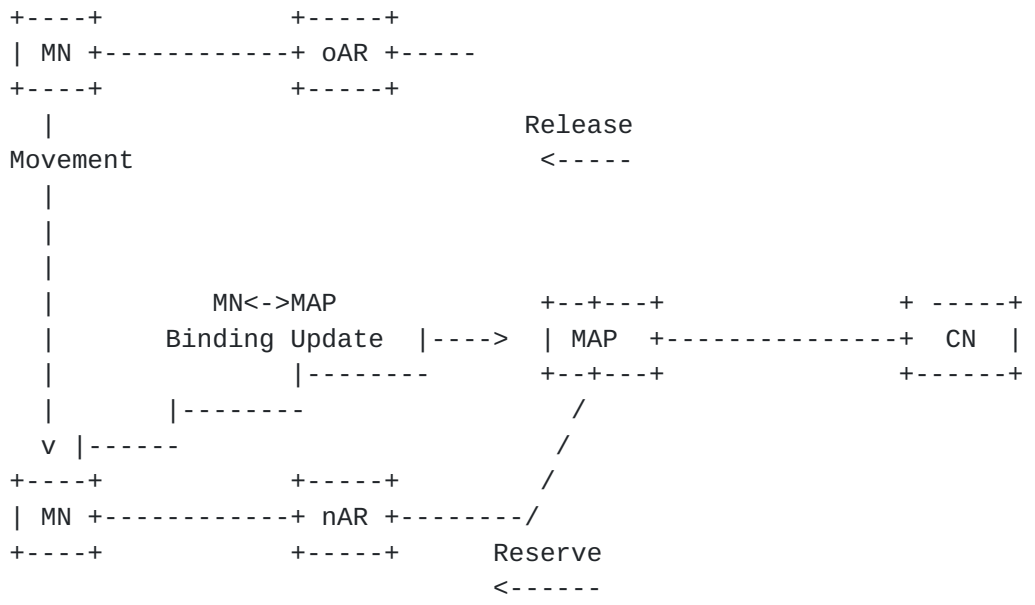

Figure 3: Downstream Reservation due to Mobility

areas:

- Session Ownership would not cause problems for mobility in the local domain since a proof of session ownership is possible by comparing an incoming request from the mobile node with the stored state at nAR in a similar manner as at oAR.

- The existing security association can be used without restarting a new authentication and key agreement protocol run. This provides performance improvements for mobility within an administrative domain. A few IPSec Context Transfer protocol proposals have already been proposed. Note that some adjustment to a security association is necessary to reflect a new care-of address.

- If a user was authorized for the indicated amount of resources then it is fair to assume that this authorization is also valid at other routers within the same domain. If the same amount of resources are available at a different path might however be a different issue. A query message could provide more information about resource availability. Avoiding an other authorization step with a possible involvement of a PDP, AAA and backend database is likely to provide performance advantages.

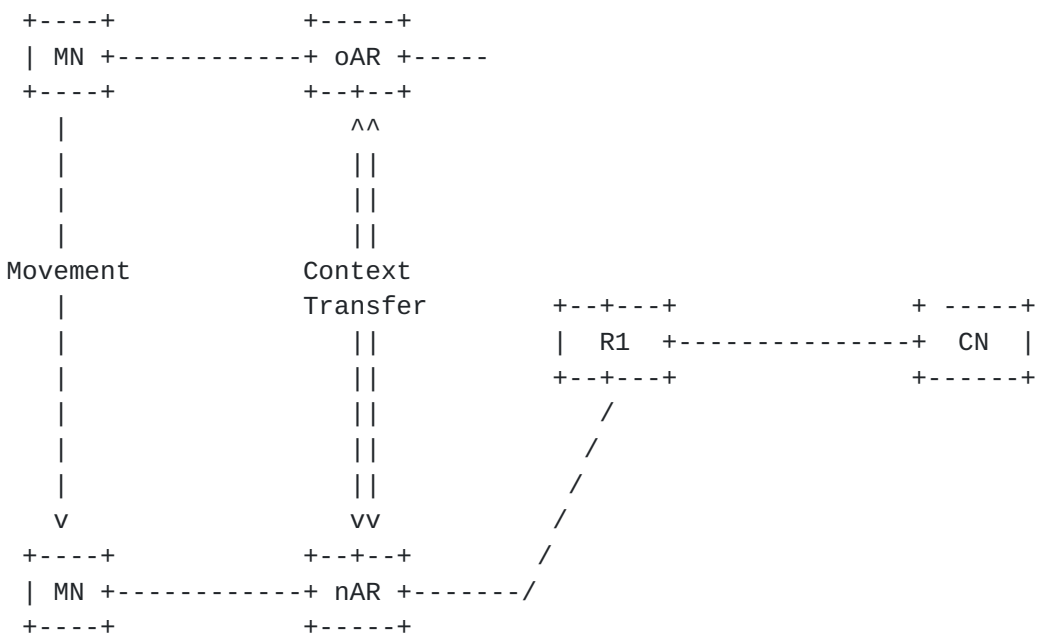

Figure 4: Context Transfer Implications

As a summary the goals of CASP (and NSIS in general) with regard to mobility signaling can be characterized as follows:

- Whenever possible the signaling message exchange caused by mobility should be kept local. This means that an end-to-end state information update along the path should not be required.
- CASP is designed to work without requiring a specific micro- and/or macro-mobility protocol. The protocol should therefore not be strongly coupled with such a protocol. An interaction with these protocols might be useful (as a trigger mechanism using an API) to provide more efficient QoS reservations.
- Mobility management introduces elements and mechanisms which modify routing. Hence it is important to have a mechanism which allows Traffic Selector updates mid-path and mid-session.
- Avoiding double reservations between the cross-over router and the correspondent node (for upstream reservations).
- Optionally removing a reservation between the cross-over router and the old access router.

Finally it should be mentioned that the selection of Traffic Selector values has some implications for signaling and security. If a regular 5-tuple is used then mobility is likely to require a signaling message exchange along the path between the MN and the CN to adjust the Traffic Selector. The same is true for a flow label and a regular end-to-end IPSec protected data traffic. If macro- and/or micro-mobility scheme is used and the Traffic Selector reflects this circumstance then a signaling message exchange can be kept local. As noted in [6] an adversary might use identity spoofing (in this case the header of ip packets) to enable its own traffic to experience preferential treatment. Marking of packets with DSCP is an extreme example which allows packet treatment without having end point addresses in the identifier.

8 Security Considerations

CASP relies on the security mechanisms described in [1]. Securing the messaging layer in a CASP peer-to-peer fashion is provided either by IPsec or 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 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 [7] and [8] can be provided without substantial protocol modification/extensions.

9 Open Issues

- CASP can use a 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.

  The Next request feature is a flexible concept. Is this degree of flexibility desired?

- CASP can use a Priority object to indicate 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.

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

reservations. 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 [9], 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.

- CASP-QOS, like CASP, can support source-specific multicast (SSM) [10]. 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.

- 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.

A CASP-QoS Message Formats

A CASP-QoS message consists of a common header, followed by a body consisting of a variable number of variable-length objects, which are identified as being of a particular type. The following subsections define the format of the common header, the standard object header, and the construction of CASP-QoS messages.

A.1 Common Header

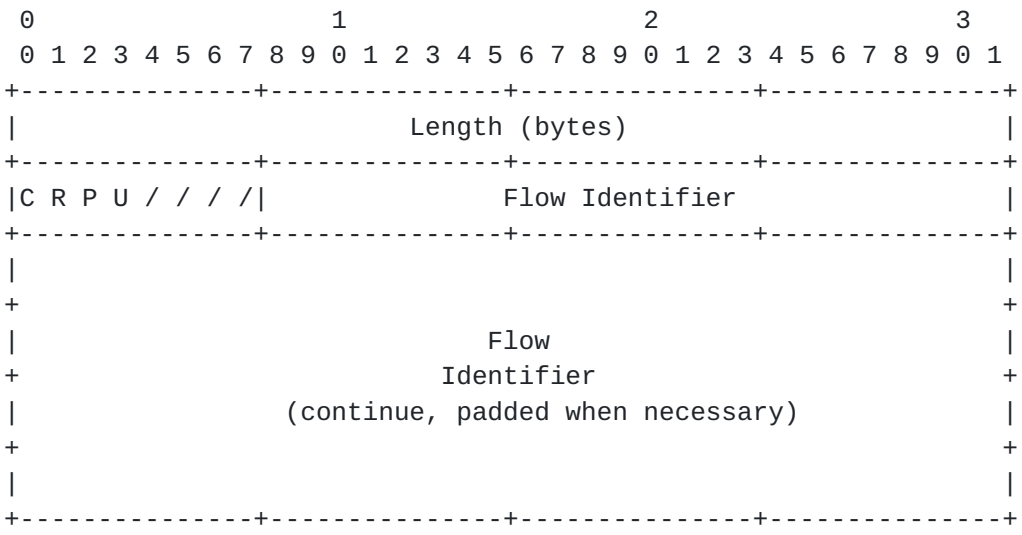

The fields in the common header are:

- Length. Indicates the length of the CASP message in bytes. This includes the length of the common header (including the length field itself).

- Flags: 8 bits

  Currently four flags are defined:

  - The Commit (C) bit indicates that the message receiver should send back a Commit message in the opposite direction.

  - The Removal (R) bit indicates that this message removes all CASP-QoS state (Reserve and Commit states, if any) for the

CASP-QoS session. If not set, the message establishes or refreshes CASP-QoS state.

- The Partial Reservation (P) bit requests that a partial reservation is acceptable. If not set, the reservation from the sender to the receiver should be tried if possible.

- The unsecure (U) (or "tainted") bit indicates that the message has traversed a hop without channel security.

· Type: 8 bits The CASP-QoS message type. Currentl valid types are:

- Type 1: CASP-QoS Reserve Message

- Type 2: CASP-QoS Commit Message

- Type 3: CASP-QoS Release Message

- Type 4: CASP-QoS Success Message

- Type 5: CASP-QoS Error Message

· Flow Identifier

Contains information about the flow which should receive a particular QoS treatment. It contains the IP addresses of the data sender and data receiver, and possibly some additional demultiplexing information (such as protocol type, source and destination ports, SPI or flow label).

A.2 Object Formats

Each object consists of one or more 32-bit words with a one word header, with the following format:

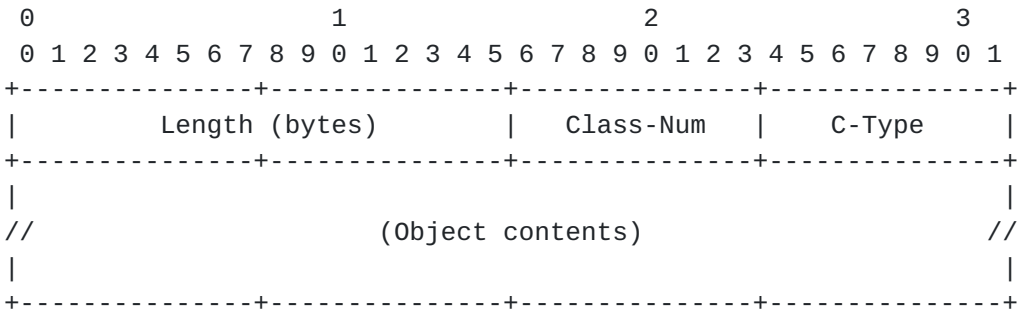

An object header has the following fields:

· Length: 16 bits

The length field contains the total object length in bytes, including the object header.

· Class-Num: 8 bits

Identifies the object class; values of this field are to be provided. Each object class has a name, which is always capitalised in this document. A CASP-QoS implementation must recognize the following classes:

- ERROR_NODE

Indicates the IPv4 or IPv6 address where a reservation failed.

- VERSION

Indicates which type of flowspec is used. This is used to easily detect a new flowspecs, rather than including the flowspec itself. Currently Version = 0 indicates IntServ Flowspec.

- C-Type: 8 bits

Object type, unique within Class-Num. Values are to be defined later.

- MINIMUM_FLOWSPEC

Indicates the minimum flowspec the application would accept.

- DESIRED_FLOWSPEC

Indicates the maximum possible flowspec the application would desire.

The format of an IntServ Controlled Load FLOWSPEC is given as follows. Note during a session, only one flowspec should be included unless the user wants nodes to "upgrade" the reservation if possible.

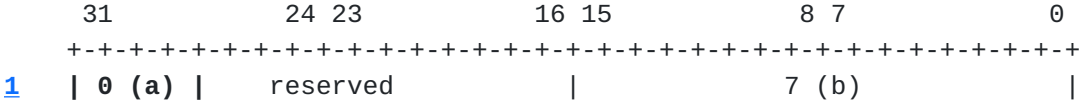

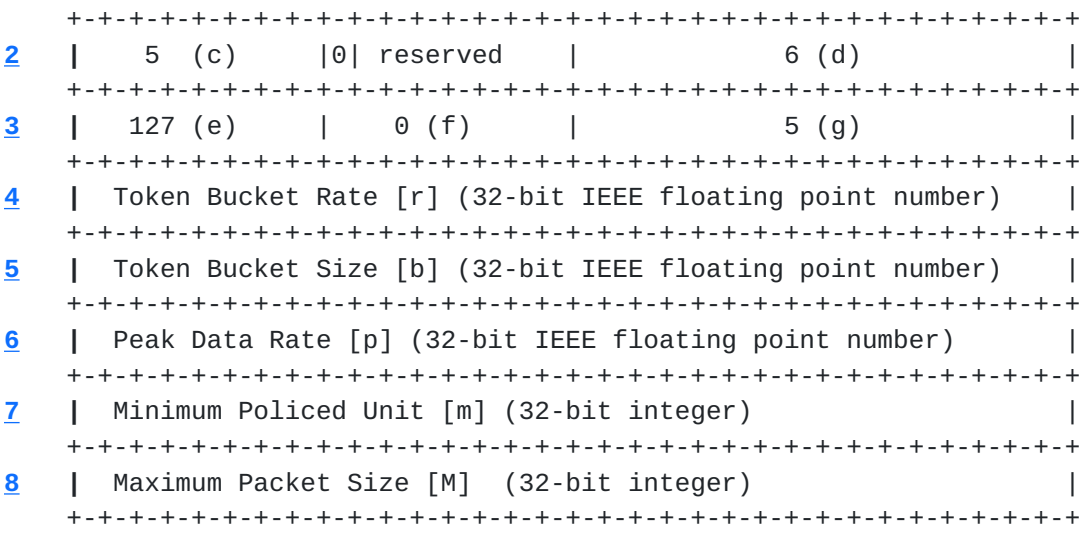

(a) - Message format version number (0)
(b) - Overall length (7 words not including header)
(c) - Service header, service number 5 (Controlled-Load)
(d) - Length of controlled-load data, 6 words not including per-service header
(e) - Parameter ID, parameter 127 (Token Bucket TSpec)
(f) - Parameter 127 flags (none set)
(g) - Parameter 127 length, 5 words not including per-service header

B 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
Institute for Informatics
University of Goettingen

Lotzestrasse 16-18
37083 Goettinge
Germany EMail: fu@cs.uni-goettingen.de

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

C Bibliography

[1] H. Schulzrinne, H. Tschofenig, X. Fu, and A. McDonald, "Casp - cross-application signaling protocol," internet draft, Internet Engineering Task Force, 2003. 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] J. Heinanen, F. Baker, W. Weiss, and J. Wroclawski, "Assured forwarding PHB group," RFC 2597, Internet Engineering Task Force, June 1999.

[5] J. Wroclawski, "The use of RSVP with IETF integrated services," RFC 2210, Internet Engineering Task Force, Sept. 1997.

[6] H. Tschofenig and D. Kroeselberg, "Security threats for nsis," internet draft, Internet Engineering Task Force, 2003. Work in progress.

[7] 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.

[8] 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.

[9] F. Dawson and D. Stenerson, "Internet calendaring and scheduling core object specification (icalendar)," RFC 2445, Internet Engineering Task Force, Nov. 1998.

[10] H. Holbrook and B. Cain, "Source-specific multicast for IP," Internet Draft, Internet Engineering Task Force, Feb. 2002. Work in progress.

Table of Contents