

Internet Draft
<draft-salsano-bgrpp-arch-00.txt>

May, 2002
Expiration: November, 2002

Stefano Salsano (ed.), CoRiTel
Vincenzo Genova, CoRiTel
Fabio Ricciato, CoRiTel
Martin Winter, Siemens AG
Bert F. Koch, Siemens AG
Lila Dimopoulou, NTUA
Eugenia Nikolouzou, NTUA
Petros D. Sampatakis, NTUA
Iakovos S. Venieris, NTUA
Gerald Eichler, T-Systems Nova GmbH

Inter-domain QoS Signaling: the BGRP Plus Architecture

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>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>.

Distribution of this memo is unlimited.

Copyright Notice

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

Abstract

DiffServ architecture provides a set of "user plane" tools for IP QoS. Deployment of DiffServ is starting with static, single-domain solutions. Dynamic QoS and multi-domain QoS are still open issues, as witnessed by the activity of the IETF NSIS WG, related to the "signaling plane". In order to realize true end-to-end QoS services in the Internet, spanning multiple administrative domains, efficient and scalable signaling and resource control mechanisms are needed.

This document describes a scalable inter-domain resource control architecture for DiffServ networks. The architecture is called BGRP

Plus, as it extends the previously proposed BGRP framework.

S. Salsano (ed.)	Expires November 2002	1
Inter-domain QoS Signaling: the BGRP Plus Architecture		May-02

Table of Contents

1. Introduction.....	4
2. Requirements for inter-domain resource reservation.....	5
3. Reference Network.....	6
4. Scalability of inter-domain resource reservation.....	7
5. Quiet grafting.....	8
6. Inter-domain QoS and Globally Well Known Services.....	10
7. BGRPP messages.....	10
8. BGRPP procedures.....	11
9. BGRP+ and BGRP.....	13
10. Acknowledgements.....	14
11. Author Information.....	14
12. References.....	15
13. APPENDIX A: PSEUDO CODE for Message Processing.....	15

1. Introduction

The DiffServ architecture specifies a set of "user plane" mechanisms, which can be used to provide QoS in IP networks. Intentionally, the IETF DiffServ WG has not covered any "signaling plane" aspect. QoS signaling capabilities are indeed needed to extend the provisioning of QoS in IP networks from a static model towards a dynamic one. The IETF WG NSIS (Next Steps In Signaling) [1] has been specifically chartered to address the signaling aspects of QoS in IP networks. The NSIS WG is currently defining the requirements for the QoS signaling mechanisms, and is considering a set of reference scenarios. One of these scenarios is the QoS reservation/ negotiation over administrative boundaries or "inter-domain" QoS.

A fundamental issue in the definition of an inter-domain QoS model is the scalability, because the ambitious goal is to support QoS services on the scale of the global Internet.

This draft proposes an architecture that originates from the BGRP protocol framework proposed in [2], providing "sink tree" based aggregation for resource reservations over a network of DiffServ domains. The aggregated reservations are negotiated between so-called BGRP agents, which are deployed at each BGP-capable border router of each DiffServ domain. In this way, each domain can perform some kind of admission control, taking into account the available resources within that domain. By aggregating the reservations according to the sink trees created by the BGP routing protocol [3], the number of reservations and thus the amount of state information stored in the network can be reduced.

However, aggregation of reservations is just the first step towards scalability. To limit the signaling load and the processing power required in the BGRP agents, it is also necessary to reduce the number of signaling messages. We propose mechanisms for the early response to reservation messages, in [2] called "quiet grafting", so that not each message has to travel edge-to-edge through the DiffServ network region. The architecture proposed in this draft is called BGRP Plus (BGRPP or BGRP+).

As regards the applicability of this proposal:

- 1) The BGRPP architecture has been implemented in the context of the AQUILA IST project, where a specific access QoS signaling mechanism has been defined [4].
- 2) The BGRPP architecture is largely compliant to the requirements for the QoS signaling under definition by NSIS WG.
- 3) Considering the IntServ over DiffServ architecture as described in [5], then an efficient and scalable resource control is also essential within the DiffServ region. BGRPP may be a more scalable answer to this need, with respect to the "RSVP-aware DiffServ region", which is proposed in [5].

S. Salsano (ed.)	Expires November 2002	3
Inter-domain QoS Signaling: the BGRP Plus Architecture		May-02

Our discussion is focused on technical aspects and even if we mention the Service Level Agreements (SLAs) the discussion of financial and legal aspects of the inter-domain services is intentionally omitted.

2. Requirements for inter-domain resource reservation

An architecture and related protocol for inter-domain resource reservation should meet the requirements listed in the following paragraphs. BGRP is a good candidate to meet them.

2.1 Scalability

The architecture has to scale well with the Internet size and growth.

2.2 Autonomous operation

Admission control should be made autonomously in each domain. Each domain's network operator can independently administer his network and configure the behavior of inter-domain resource reservation elements independently each other.

2.3 Follow established routes

Inter-domain resource reservation shall rely on inter-domain routing (e.g. BGP) for routing decisions. An interface has to be defined between inter-domain routing and inter-domain resource reservation to allow the latter to retrieve routing information (e.g. the NEXT_HOP). We assume that BGP remains unaffected by the inter-domain resource reservation architecture.

2.4 Independence of each domain

The inter-domain protocol should be independent of the intra-domain resource control architecture. A domain may use static provisioning as well as dynamic resource allocation. However, a common interface between intra-domain and inter-domain resource control has to be defined.

2.5 Influence of SLA and SLS

In the inter-domain scope, the most important constraint is not an optimization of network resources but the conformity to the SLA with the neighbor. Each domain has many neighbors, from just a few up to hundreds, and the administrator signs a contract with each of them. When a request for a reservation to another AS reaches a BR, it has to choose a "next hop" to forward the request: if the next hop belongs to another AS, the BR should make sure that the request is in compliance with the SLA.

S. Salsano (ed.) Expires November 2002 **4**
 Inter-domain QoS Signaling: the BGRP Plus Architecture May-02

3. Reference Network

The architecture described in this document assumes a DiffServ region consisting of several connected, but administratively separated domains. The traffic can enter and leave the domains at two different types of routers:

- An edge router (ER) connects a domain to a network, which is not taking part in the BGRPP resource allocation mechanism, e.g. an access network.
- A border router (BR) connects a domain to another domain, which also takes part in the BGRPP resource allocation mechanism.




Figure 1: Reference network configuration

A source or destination domain is a domain with at least one edge router. A transit domain is a domain with at least two border routers, which forwards traffic received at one border router to another border router. All edge routers and border routers are required to run BGP.

Corresponding to each border router, a BGRP agent is instantiated as shown in the following figure:




Figure 2: BGRPP agents

We assume that the source domain is capable of performing some kind of per flow admission control, taking into account the available resources up to BR1. The source domain has however no information about the availability of resources along the further path of the reservation through other domains.

To perform inter-domain admission control, the source domain determines the egress border router and contacts the corresponding BGRPP agent. Through the BGRPP protocol this agent is able to determine, whether resources are available in each domain along the path and on the links connecting the domains. Each BGRPP agent cares for the resources on the next path segment from its associated border router towards the destination.

Resource reservation in the source and destination domain is not the task of the BGRPP protocol. However, as we describe later, BGRPP has to provide mechanisms to enable resource reservation in the destination

domain.

There are mainly two types of messages involved in the reservation set-up:

- PROBE messages are initiated at the source domain and are forwarded between BGRPP agents along the BGP path. They check the conformance of the request with the SLAs between neighbored domains. The path is recorded in the message, to enable the response to take the same way in the backward direction.
- GRAFT messages indicate the availability of resources towards the destination domain. During message processing, the resource availability is checked and the resources are actually reserved in each path segment.

4. Scalability of inter-domain resource reservation

BGRPP is an inter-domain protocol to allow resource reservation. The main issue that BGRPP should tackle is the scalability problem. This problem is related to two aspects:

- the handling of state information for each reserved flow and
- the processing of reservation messages.

In particular, the scalability problem is related to:

- memory needs for each router that runs reservation protocol;
- CPU usage for message processing;
- bandwidth utilization for messages.

A scalable solution will minimize these three elements.

BGRP as described in [2] mainly addresses the scalability in terms of the amount of state information kept in each BGRP agent. It aggregates

S. Salsano (ed.) Expires November 2002 6
Inter-domain QoS Signaling: the BGRP Plus Architecture May-02

reservations along the BGP sink trees and thus achieves a scalability behavior, where the memory requirements are proportional to the number of sink trees with simultaneous active reservations.

However, in order to check the resource availability, BGRP messages have still to travel along the full path to the destination domain, asking each BGRP agent to check the resource availability for that part of the network it is responsible for. So the message processing overhead may be rather large, especially in big backbone networks.

In order to solve this problem, an hint is given in [2]: resources may

be kept in advance at a BGRP agent, so that further requests may be already terminated at an earlier stage. In this document we provide the complete functional specification of this mechanism, called "quiet grafting". Together with the sink-tree-based aggregations, this mechanism provides a scalable solution for inter-domain resource reservations.

5. Quiet grafting

5.1 Requirements

When an inter-domain reservation is initiated at a source domain, the first BGRPP agent constructs a PROBE message indicating the amount of bandwidth required and the destination address. It is not evident, to which sink tree this reservation will belong. So the PROBE message is forwarded hop-by-hop between BGRPP agents along the BGP route to the destination. Obviously, the last BGRPP agent, which corresponds to the root of the sink tree, can assign a sink tree id to the reservation. The GRAFT message sent back will contain this sink tree id. As this message travels back to the source domain, it installs the necessary sink tree reservations in the path segments.

To enable an intermediate BGRPP agent to answer a PROBE message successfully with a positive response, the following conditions have to be met:

- The BGRPP agent must be able to determine the sink tree, to which the reservation belongs.
- The BGRPP agent must have pre-reserved resources for this sink tree, so that he can guarantee, that the resources are available on the path segment from the current point to the destination domain.
- As the last BGRPP agent may no longer be informed about a new reservation, the BGRPP agent must provide means to contact the destination domain, so that resources can also be reserved on the not-BGRPP-controlled path segment from BR4 to ED2 (see figure 2).

The following paragraphs describe the solutions we propose to solve these requirements.

5.2 NLRI Labeling for Sink Tree Identification

According to [2], a sink tree is identified by the destination AS number and a border router id. The network layer reachability information

(NLRI) associated with the root of the tree can be used to identify the tree at a distant point in the network. As BGP may aggregate this information into aggregated routes, it is not directly derivable from BGP routing information. Instead, we propose to propagate the NLRI back from the root of the sink tree in GRAFT messages and to store this information in each BGRPP agent, which processes the message.

When a new PROBE message arrives, the BGRPP agent will compare the destination address with the already stored NLRI thus trying to identify the sink tree. Successful sink tree identification is a prerequisite to the following steps, to perform successful quiet grafting.

To reduce memory requirements, BGRPP agents will only store NLRI information for those sink trees, where actual reservations exist.

5.3 Pre-reservation

When a BGRPP agent processing a PROBE message was able to determine the sink tree, he will now check, whether he has pre-reserved resources on this sink tree. If this is the case, he can generate a GRAFT message and terminate the PROBE at this stage. If there are not sufficient resources available, the BGRPP agent will further forward the GRAFT message to the next BGRPP agent, as determined by the NEXT_HOP attribute of the BGP path to the destination.

The amount of pre-reserved resources that is available at a given time for a sink tree is called resource cushion. A possible solution to build this resource cushion at each BGRPP agent is the "delayed resource release" mechanism. This means, that a BGRPP agent will not immediately release unused resources, but instead keep them in an attempt to satisfy further requests. Resources are however released, when they are unused for some time.

We will describe the proposed algorithm in a separate document and provide information about its performance in various scenarios.

5.4 Signaling in the last domain

When a PROBE message is answered "in advance", the last domain may not be informed about the new request and cannot reserve resources within that domain. To include the resource availability in the last domain in the overall admission control, we propose to back-propagate a reference to the interface to intra-domain resource control at the root of the sink tree within the GRAFT message. This information is also stored within each BGRPP agent that handles this message.

The originating BGRPP agent can then use this reference to directly request the required resources in the destination domain.

This reference is also necessary for the release of a reservation: since resources are not immediately released, when the original end-user request is removed, possibly the destination domain will not be informed at all about this event. For this purpose it is necessary that the originator of the request can directly inform the destination domain.

6. Inter-domain QoS and Globally Well Known Services

In the previous sections we have mentioned "reservations" which are originated by a source domain (and propagated up to the destination domain) and are characterized by a bandwidth parameter. Actually, a reservation should be also be characterized by the required service and other parameters besides the bandwidth could be needed. We assume that a set of "Globally Well Known Services" is defined to characterize the required service. For the sake of simplicity, in the following description of messages and procedures, we also assume that the bandwidth is the only needed parameter.

7. BGRPP messages

7.1 PROBE message

It consists of a reservation request and destination network information. It travels from origin AS towards destination AS and collects routing information. It contains the information of the requested amount of resources.

The fields of PROBE message are:

- Sender: It is a BGRPP agent identifier, it is composed of AS id and the IP address of BR that has sent the PROBE.
- ProbeId: The origin AS chooses the ProbeId. The ProbeID scope is local to the originating BGRPP agent and only used there to match the GRAFT (or ERROR) message. In other words, only the originating AS stores a "PROBE state" while Intermediate ASs nodes does not need correlate PROBE and GRAFT.
- GwksId: GWKS id
- Destination: IP address prefix of the host or network that is destination of request
- RequiredBW: Requested bw [bit/s]
- TreeId: Id of sink tree, it is NULL if the sink tree is unknown. This information is actually redundant, as each node could check weather the Destination matches an existing sink tree. By avoiding this check, it enhances the performances.
- Path: Record of the route in terms of BGRPP agents from origin AS

7.2 GRAFT message

When the destination AS (or a transit AS if it can do quiet grafting) receives a PROBE and it can accept the request, it returns a GRAFT

S. Salsano (ed.)

Expires November 2002

9

Inter-domain QoS Signaling: the BGRP Plus Architecture May-02

towards the origin AS, using the Path information collected by PROBE message. This message carries sink tree information. It also carries the amount of resources that are reserved by the node sending the GRAFT.

The fields of GRAFT message are:

Sender: It is a BGRPP agent identifier, it is composed of AS id
and the IP address of BR that has sent the GRAFT.
Id: Msg id to match GRAFT with PROBE, echoed from PROBE
GwksId: GWKS id
Destination: Ip address prefix of the host or network that is
destination of request, echoed from PROBE
ReservedBW: Reserved bw [bit/s]
TreeId: Id of sink tree
DestResMgr: Reference to destination domain reservation manager for
reservation in the last domain
AddressPrefixList: NLRI
Path: Record of the route in terms of BGRPP agents to be
followed back

7.3 REFRESH message

It contains the indication of the current amount of needed BW.
It is sent by a previous (upstream) node to reduce the amount of requested resource. A REFRESH with zero bandwidth is used to tear down a reservation.
REFRESH messages must never be used by a previous (upstream) node to increase the amount of requested resources.

Sender: It is a BGRPP agent identifier, it is composed of AS id
and the IP address of BR that has sent the REFRESH.
GwksId: GWKS id
ActualBw: actual reserved bandwidth
TreeId: id of sink tree

7.4 ERROR message

It indicates that some error has occurred. It contains a description of the specific error case. For example if the resources are not available, an ERROR message is sent and propagated backwards up to the originator of the request.

7.5 TEAR message

It was defined in the BGRP architecture. Currently, it has no use in the

BGRPP architecture, because it is replaced by the use of refresh messages.

8. BGRPP procedures

8.1 State information in BGRPP agents

In order to process the BGRPP protocol messages, the BGRPP agent stores the following sink tree status information:

For each sink tree (with at least one reservation) and for each service, the next ("outgoing") hop and a list of previous ("incoming") hops is stored. The value of reserved resources for the outgoing hop and for each incoming hop is stored (this information is replicated per each GWKS).

For each sink tree, the NLRI is stored, to enable sink tree identification for quiet grafting (see 5.2). Additionally, a reference to the intra-domain resource control of the destination domain is stored, to enable signaling to the last domain (see 5.4).

TABLE 1 - Sink tree x state information

Next hop	IP address of next hop BGRPP agent
Outgoing.res[g]	Reserved BW for the sink tree x towards the Next hop for the GWKS g
Previous hop[i]	IP addresses of previous hop BGRPP agents
Incoming[i].res[g]	Reserved BW for the sing tree assigned to each Previous hop for the GWKS g
NLRI	NLRI for the sink tree
Intra-dom. RC Ref.	Reference to the intra-domain resource control of the destination domain

The resource cushion is the difference between the Outgoing.res and the sum of the Incoming[i].res.

8.2 Message handling

(In the following the dependence from the GWKS g will always be omitted).
The message handling for a transit AS is described, the procedures for Originating AS and Terminating AS can be derived. In APPENDIX A a more complete pseudo-code description of the message handling is provided.

- PROBE messages

The PROBE messages can be:

- rejected because SLA/SLS does not match or there are no resources on outgoing link or on intra-domain links
- accepted with Quiet Grafting
- forwarded to downstream node with no change in RequiredBW

When a PROBE is received a preliminary admission control is performed. If the BGRPP agent is at an ingress Border Router, it checks for SLA between the requester AS and its own AS and decides if the request can be accepted. The BGRPP agent could also check if intra-domain resource manager could accept the new request. If the BGRPP agent is associated to an egress BR it has to check the SLA with the AS of next hop BR. It can also check for the resources on the outgoing link towards the next hop BR.

S. Salsano (ed.) **Expires November 2002** **11**
Inter-domain QoS Signaling: the BGRP Plus Architecture May-02

If any of these checks fail, the request cannot be accepted and an ERROR message is sent to previous hop.

If the checks are OK, the node (both the ingress and the egress BR) first verifies if the Destination of the PROBE matches an existing sink tree. In this case it is first checked if the resource cushion can fit the RequiredBW. In case of positive answer, Quiet Grafting is performed: the Incoming[PH].res related to the previous hop is increased by the RequiredBW value and a GRAFT message is sent to the Previous hop. If the Destination of the PROBE does not match an existing sink tree, or the resource cushion is not enough, the PROBE is forwarded and the state information of the node is not touched in any way.

- GRAFT messages

The resource reservation is performed when a GRAFT arrives. A GRAFT tells the receiving node the downstream node can accept a given amount of BW for a given sink tree. Now the receiving node should decide if it has the resources to send this amount of BW towards the next hop BR and if the SLA matches. In particular, an Ingress BR should check the Incoming SLS with upstream AS and should consider the availability of intra-domain resources towards egress BR. An Egress BR should check the Outgoing SLS towards the downstream AS and should consider the resources on the outgoing link towards the Ingress BR of the downstream AS. (The nodes could have already performed these checks during PROBE phase, but the situation could have changed).

If the resources are available and SLA/SLS matches, the graft is accepted and the node:

- 1) Increases the Outgoing.res by the ReservedBW value.

2) Increases the Incoming[PH].res of the Previous hop PH by the ReservedBW value.

3) Propagates the GRAFT to the Previous hop.

If the resources are not available or (incoming/outgoing) SLA/SLS not matching, the GRAFT cannot be accepted. The node will immediately send an ERROR towards the first originator of the PROBE. If the reason of refusal is outgoing SLA not matching or unavailability of resources, the node has also to release reserved BW to next hop because it can not utilize it. To this purpose the node sends a REFRESH message to its next hop. If the reason of refusal is incoming SLA not matching the node can maintain the additional reserved BW towards its next hop to enlarge resource cushion.

9. BGRP+ and BGRP

It is worth to recall the extensions of this draft with respect to the original BGRP proposal:

- The architecture includes the aspects related to the interactions with intra-domain resource reservation mechanisms
- The quiet grafting mechanisms simply mentioned in the BGRP proposal has been fully specified
- The mechanisms to match a flow into the sink-tree by means of the NLRI information and to distribute this NLRI information where needed have been specified

S. Salsano (ed.)	Expires November 2002	12
Inter-domain QoS Signaling: the BGRP Plus Architecture		May-02

- The issue of the missing signaling in the last domain due to quiet grafting has been addressed and resolved
- The semantic content of the messages has been described
- The procedures for message handling have been given (both in text and using a pseudo-code)

10. Acknowledgements

Part of this work has been funded under the European Commission 5th framework IST program.

The authors would like to acknowledge all their colleagues in the AQUILA project for their important contribution to this work.

11. Author Information

Stefano Salsano
CoRiTel - Consorzio di Ricerca sulle Telecomunicazioni
Via Anagnina, 203

00040 Morena Roma - ITALY

e-mail: salsano@coritel.it

Vincenzo Genova, CoRiTel

e-mail: genova@coritel.it

Fabio Ricciato, CoRiTel

e-mail: ricciato@coritel.it

Martin Winter,

Siemens AG, ICN WN CC EK A19

81359 Munich - Germany

e-mail: martin.winter@icn.siemens.de

Bert F. Koch, Siemens AG

e-mail: bert.koch@icn.siemens.de

Lila Dimopoulou

NTUA - National Technical University of Athens

Heroon Polytechniou 9, 157 73, Athens GREECE

e-mail: lila@telecom.ntua.gr

Eugenia Nikolouzou, NTUA

e-mail: enik@telecom.ntua.gr

Petros D. Sampatakis, NTUA

e-mail: psampa@telecom.ntua.gr

Iakovos S. Venieris, NTUA

e-mail: ivenieri@cc.ece.ntua.gr

Gerald Eichler, T-Systems Nova GmbH

S. Salsano (ed.)

Expires November 2002

13

Inter-domain QoS Signaling: the BGRP Plus Architecture

May-02

e-mail: Gerald.Eichler@t-systems.com

12. References

- [1]M. Brunner (Editor), "Requirements for QoS Signaling Protocols",
[draft-ietf-nsis-req-01.txt](#), April 2002, Work in Progress
- [2]P. Pan, E. Hahne, H. Schulzrinne: "BGRP: Sink-Tree-Based
Aggregation for Inter-Domain Reservations", Journal of
Communications and Networks, Vol. 2, No. 2, June 2000, pp. 157-167,

<http://www.cs.columbia.edu/~pingpan/papers/bgrp.pdf>

- [3]Y. Rekhter, T.J. Watson, T. Li: "A Border Gateway Protocol 4 (BGP-

- 4)", [RFC 1771](#), March 1995
- [4]AQUILA IST Project <http://www.ist-aquila.org/>
- [5]Y. Bernet et al., "Integrated Services Over Diffserv Networks", [RFC 2998](#), November 2000

13. APPENDIX A: PSEUDO CODE for Message Processing

13.1 PROBE processing

```
// State variables for PROBE processing:
treeOutgoing.res
treeIncoming[i].res

// information related to messages:
probe.required
new_graft.reserved
new_probe.required
new_error.failedBW

// Local variables:
cushion
request

// PROBE processing:

// perform BGP lookup to determine the BgrpNeighbour

// test incoming SLA
if (no resources in SLA)
    new_error.failedBW=probe.requiredBW
    send error message to previous hop
    exit
endif

// try to identify the sink tree (DestinationDomain) for the message
if (PROBE does not contain a tree id)
    // perform lookup in NLRI table to find the tree
    if (tree found)
        insert tree id in PROBE
    endif //
endif
```

```
//check quiet grafting possibility
request = probe.required
cushion=treeOutgoing.res-sum(treeIncoming[i].res)
if (cushion>=request && tree found)
```



```

    // perform quiet grafting
    treeIncoming[IN].res += request
    // send graft message to previous hop
    new_graft.reserved=request
else
    // quiet grafting is not possible
    // test outgoing SLA
    if (no resources in SLA)
        // send error message to previous hop
        new_error.failedBW=probe.required
    else
        // prepare and forward PROBE message to next hop
        // update route hop
        new_probe.required=request
    endif
endif
endif

```

13.2 GRAFT processing

```

// State variables:
treeOutgoing.res
treeIncoming[IN].res

// information related to messages:
graft.reserved
new_graft.reserved
new_error.failedBW

// GRAFT processing:

// check outgoing SLA and at BGRPP ingress node request intradomain
resource
if (no resources available)
    // send error message to previous hop
    new_error.failedBW= graft.reserved
    // send refresh message to next hop
    new_refresh.actualBW=treeOutgoing.res
else
    // update reserved bandwidth
    treeOutgoing.res += graft.reserved
endif

// check incoming SLA
if (no resources available)
    // send error message to previous hop
    new_error.failedBW= graft.reserved
else
    // update reserved bandwidth

```

```
    treeIncoming[IN].res += graft.reserved
    // prepare and forward GRAFT message to previous hop
    update route hop
    new_graft.reserved = graft.reserved
endif
```

13.3 REFRESH processing

```
// State variables:
treeIncoming[IN].res
treeOutgoing.res

// information related to messages:
refresh.actualBW
refresh.sender
new_refresh.actualBW

// REFRESH processing:

if(refresh.sender is next hop for tree)
    // next hop knows more than I
    if (refresh.actualBW > treeOutgoing.res)
        error situation
        //Send to next hop a REFRESH message
        new_refresh.actualBW=treeOutgoing.res

    // next hop thinks less than I
    if (refresh.actualBW < treeOutgoing.res)
        error situation
else
    // previous hop wants less than I know (like a TEAR)
    if (refresh.actualBW < treeIncoming[IN].res)
        treeIncoming.res = refresh.bw
    // previous hop thinks more than I
    if (refresh.actualBW > treeIncoming.res)
        error situation
end if
```

13.4 ERROR processing

```
// information related to messages
error.failedBW
new_error.failedBW

// ERROR processing
if(this hop is destination)
    advise intradomain requester
else
    // prepare and forward ERROR message to previous hop
    // update route hop
```

```
new_error.failedBW = error.failedBW
```

[S. Salsano](#) (ed.)

Expires November 2002

16