566

Evaluation of a Tunnel Set-up Mechanism in QoS-aware Ethernet Access Networks

Filip De Greve, Frederic Van Quickenborne,
Pim Van Heuven, Filip De Turck, Brecht Vermeulen, Steven Van den Berghe, Ingrid Moerman, Piet Demeester, Sven Van den Bosch*, Nico Janssens*, Christele Bouchat*, Bert Van Vlerken*

Department of Information Technology, Ghent University, St-Pietersnieuwstraat 41, B-9000 Ghent
*Alcatel Bell, Research and Innovation, Francis Wellesplein 1, B-2018 Antwerp

E-mail: filip.degreve@intec.ugent.be, frederic.vanquickenborne@intec.ugent.be

Abstract **This paper describes the design of a novel tunnel set-up mechanism for Ethernet access networks. The establishment of VLAN-based tunnels is realised with a "Scoped Refresh" extension of the GVRP standard. For the distribution of QoS-related reservation parameters, a new GARP-based protocol (G2RP) is developed with closed loop design, a mechanism for updating reservation parameters and hop-by-hop admission control. Both protocols are implemented on a Linux test bed using the Click Modular Router. Extensive performance evaluations were performed. The results are compared with Linux implementations of other tunnel set-up mechanisms.**

Keywords **Ethernet, tunnel, GARP, GVRP, G2RP, Click**

I. Introduction

A. Motivation

Service Providers (SPs) are offering a wide range of services to their subscribers such as managed home networking, multimedia content delivery, video phoning, on-line gaming etc. A QoS-aware access network is necessary for connecting the SPs and the end users. The typical structure of such an access network is represented in Fig. 1. The Access Gateways (AGs) are situated relatively close to the subscribers and are concentrating groups of end users. The Service Gateways (SGs) provide the connectivity with the SPs' networks. The Access Network Providers (ANPs) would like an

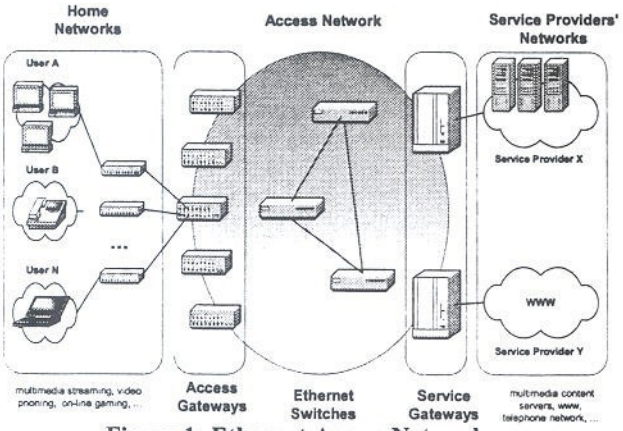

Figure 1: Ethernet Access Network

efficient way of maintaining the access network under all (dynamically changing) circumstances. Simply overdimensioning the link bandwidths to fulfill all QoS-related requirements is not the optimal solution that ANPs prefer. Subscribers should be able to exploit the offered services (or newly added services) in a dynamic way, with low access times and with high availabilities. This requires not only a fast and flexible tunnel set-up mechanism to provide the connectivity but also an associated admission control to maintain QoS guarantees for the admitted connections.

Mainly for economical reasons ANPs [1] tend towards access networks consisting of standard QoS-aware Ethernet switches (IEEE 802.1d, IEEE 802.1q & p compliant). However, the existing tunnel mechanisms for MPLS/IP networks (such as LDP, RSVP and their extensions) are not portable to Switched Ethernet networks. Explicit tunnel mechanisms are not possible because Ethernet switches are not explicitly aware of their neighboring switches. Also, tunnel mechanisms making use of IP routing tables are not adequate because the entries of the Forwarding Databases are regularly deleted: explicitly flushed or timed out. It is obvious that these problems can only be resolved by a protocol with broadcasting capabilities.

B. Contribution

For establishment of the tunnels, we specifically developed a "Scoped Refresh" extension of the standardized GVRP. This extension resolves the relevant problem of depletion of the limited sized VLAN Databases (detailed in section II). For distributing the reservation parameters associated with these tunnels, a new GARP application, referred to as G2RP, is developed with its own specific GIP context, closed loop design, mechanism for updating reservation parameters and hop-by-hop admission control. Both protocols are implemented on a Linux test bed using the Click Modular Router and extensive performance evaluations were performed.

C. Related Work

The Generic Attribute Registration Protocol (GARP), described in IEEE 802.1d [2], provides a generic attribute dissemination capability for bridged LANs. There exist two standardized GARP applications (described in IEEE 802.1q [3]): GARP Multicast Registration Protocol (GMRP) and GARP VLAN Registration Protocol (GVRP). GMRP allows group membership information to be distributed but we are specifically interested in GVRP, which provides a mechanism for dynamic maintenance of the VLAN Active Filtering Databases and for propagating the information they contain to other VLAN-aware bridges. Commercial protocol developers such as [4] have created GVRP implementations and

most Ethernet switch vendors support GVRP but to our current knowledge there is no open source implementation of a GVRP daemon for Linux available. The GARP timer behaviour has been extensively studied in [5] for a shared LAN environment but it is no goal of this paper to extend this study.

II. GVRP Extension And G2RP Design

A. GVRP Extension

The choice for VLAN-based tunnels (e.g. instead of MPLS) is quite straightforward, knowing that most of the (cheap) Ethernet switches already support VLANs. In the search for a new tunnel set-up mechanism, we found that GARP-based protocols are providing all the necessary capabilities. Therefore, why develop a brand new protocol, if you can intelligently make use of a standardized Ethernet broadcasting protocol? GVRP removes the burden of manually installing and maintaining VLANs from the network administrator's hands. The automatic VLAN registration is performed in a more consistent and reliable way compared to the laborious manual VLAN configuration on every switch in the network. GVRP not only reduces the chance of incorrect VLAN configurations but also makes the VLANs resilient to Layer 2 network failures because it works in conjunction with the spanning tree protocol (STP). After the spanning tree protocol has converged, the VLANs are automatically remapped to the new topology induced by the new spanning tree (ST). However, standard GVRP suffers from one big drawback.

Consider the following access network: the network contains a single SG and 2000 AGs. The network devices are Ethernet switches supporting GVRP with VLAN Databases of 64 entries (most currently available switches do not support the whole range of VLANs) and one Ethernet switch with a VLAN Database of 2000 entries to connect to the SG. The exact amount of switches and the chosen topology are not relevant for the considered problem but in this legitimate situation, it is impossible to set up an individual VLAN pipe for every SG-AG pair with standard GVRP-based methods. GVRP distributes all the registered VLAN IDs to every switch in the network and produces a lot of overhead registrations that will flood the limited sized VLAN Databases. GVRP was namely never designed to set up point-to-point VLAN tunnels but the VLAN topologies would rather resemble a sub-tree of the ST topology. This VLAN Database depletion problem can be resolved by removing the overhead registrations and therefore we developed the "Scoped Refresh" extension of GVRP. In Fig. 2, the GVRP operation to set up a tunnel between switches A and B is illustrated. Just like RSVP and LDP, a single GVRP run is responsible for setting up a single direction of the tunnel, meaning that a trigger to each end device is required to configure the entire tunnel. First, a single edge device is triggered and the VLAN ID is propagated towards the entire network (Fig. 2a). After triggering the other end device, the bi-directional VLAN path between A and B is established but the overhead registrations remain (Fig. 2b). For the GVRP extension, the desired end situation is shown in Fig. 2c. Every switch will detect individually the number of registrations for a specific VLAN value on its interfaces. If this number equals two, the switch is part of the VLAN tunnel and will issue de-registration messages on its non-registered interfaces for this VLAN value. If this number becomes lower than two, the switch is no longer part of the tunnel and will issue registration messages on its non-registered interfaces.

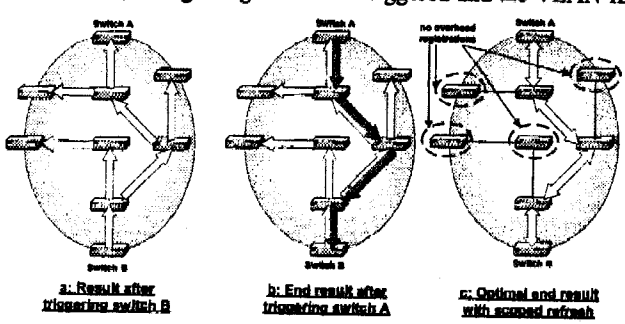

Figure 2: Scoped refresh extension

B. G2RP Design

The extended GVRP resolved setting up a physical point-to-point tunnel but is still lacking means to configure the switch hardware according to the QoS-related parameters associated with each tunnel. Instead of extending GVRP to support attribute types to propagate reservation parameters, a new GARP protocol is designed. This protocol is called "GARP Reservation Parameters Registration Protocol" (G2RP). By separating the distribution of reservation parameters and VLAN IDs, G2RP remains independent of the applied tunnel set-up mechanism (e.g. static VLAN entries).

G2RP is designed according to the GARP standard. We will address the unresolved issues of the standard that are implementation specific. First of all, the G2RP application address needs to be defined. This address will be used as destination address in the Layer 2 header of the G2RP PDUs. It is chosen from the range of GARP broadcast addresses: 01-80-C2-00-00-20 to 00-80-C2-00-00-2F. The first two addresses from the range are in use by the two existing GARP applications GMRP and GVRP [3]. We have chosen the first free address in the list: 01-80-C2-00-00-22.

Secondly, the G2RP specific attribute type and values need to be defined. Only a single attribute type is required and the attribute value format is represented on Fig. 3. The first field indicates the

VLAN-ID	RSV-ID	Reservation Parameters A	Reservation Parameters B

Figure 3: G2RP Attribute Values Format

VLAN associated with the reservation parameters and the second field is a reservation session ID, followed by the two fields of reservation parameters, one field for each direction of the tunnel. These reservation parameters contain information such as the QoS class, bandwidth parameters, burst size and the size of the time sample window. G2RP allows a robust method for making reservations per VLAN and per QoS class. Existing reservations can be altered by distributing updated parameters for the same VLAN value and QoS class. A different RSV-ID is used to distinguish incoming PDUs with updated parameters from incoming PDUs due to the GARP refresh mechanism.

After defining the G2RP semantics, the appropriate hardware operations have to be linked with registration and de-registration events. Basically, the switch hardware (bandwidth shapers, classifiers, queues, etc.) will be configured according to the registered parameters for the associated VLAN value and QoS class. Before any configuration takes place, G2RP will consult the admission control. Admission control is added to keep track of the existing bandwidth reservations. If at any point along the VLAN tunnel the available hardware resources are not sufficient to support a specific reservation, this reservation will fail and its PDUs won't be propagated in the switch.

Finally, the GIP (= GARP Information Propagation) context is free to define. This context defines a subset of the ports of the switch,

forming the active topology for the propagation of GARP messages. GMRP and GVRP are both using the active ST topology as GIP context. The GARP standard allows more complex active topologies depending of the instance of ST or VLAN ID. For the G2RP application, we linked the GIP context to the pre-established VLAN tunnel. This allows safe broadcasting of reservation parameters in the associated VLAN tunnel. While the GIP context was previously fixed for all GARP applications, G2RP introduces a dynamic GIP context. For every incoming G2RP PDU the associated VLAN ID is identified and used for the propagation towards outgoing ports, which are member of the same VLAN.

Just like GVRP, a single run of G2RP is responsible for configuring a single direction of the tunnel but in contrast to the GVRP, a G2RP reservation session can fail due to interference of the admission control. Therefore, a closed loop design is preferred for G2RP. This is illustrated in Fig. 4. In contrast with GVRP, a single trigger is sufficient to configure the entire tunnel: at the opposite end of the VLAN tunnel (i.e. switch B), the G2RP PDUs are reversed and sent back to the originating end of the tunnel (i.e. switch A). G2RP PDUs will always follow the VLAN pipe, even in

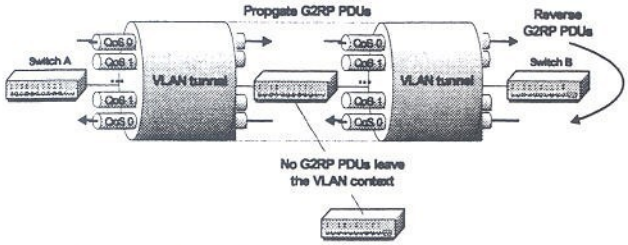

Figure 4: G2RP message propagation

conditions of VLAN path reconfiguration. If at the originating switch the reversed PDUs do not arrive in time, the reservation session will have failed due to insufficient network resources.

III. Implementation Aspects

The extended GVRP and the novel G2RP application are implemented by means of a platform independent C code and both applications are tested on a Linux test bed (see Fig. 5). The GVRP/G2RP daemon is running on every Linux PC of the network. The interaction between the end user and the daemon is realized via a dedicated CLI (Command Line Interface) socket. By using standard sockets it is possible to send requests from a single host (or management system) by means of the CLI to all the daemons. The Data Plane is emulated using the Click Modular Router Toolkit.

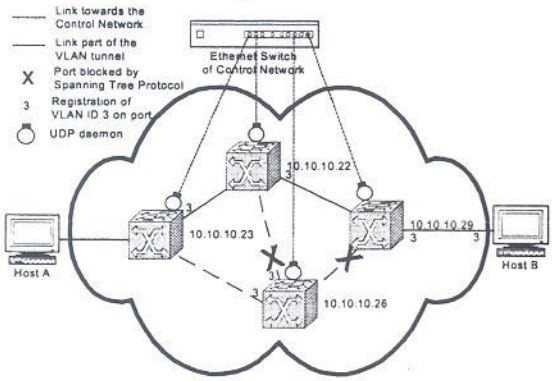

Figure 5: Linux Test Bed

Click is a software architecture for building flexible and configurable routers. Using newly developed elements in combination with standard Click elements, the Click Router is configured as a fully VLAN compliant Ethernet switch. The platform independent code is extended with an Adaptation Layer containing Linux-related code for sending Ethernet frames towards the physical interfaces, code for receiving Ethernet frames from the Data Plane and the remaining code supports the interaction with Click: editing the VLAN Filter Databases, modifying the configurable switching hardware and accessing the STP port states. For further details concerning the Click Router, the protocol development or other implementation aspects, we refer to the Linux Kongress paper [6].

IV. Evaluation Results

We tested two scenarios: service realization and failure recovery on the Linux test bed (Fig. 5). We also tested the influence of logging on the protocol performance and compared the measured results with a Linux implementation of RSVP-TE.

A. Scenario 1: Service Realisation

If a new point-to-point service gets activated, VLAN tunnels have to be realised between every AG and the SG of the Service Provider. The SG will be triggered to register a VLAN for every AG in the network. On their turn, the AGs can complete the tunnel set-up by dynamically registering VLANs if needed. On our test bed we measured the time it takes to perform a burst of sequential tunnel set-ups. The set-up times are measured at one end device of the tunnel (i.e. 10.10.10.23, see Fig. 5). This end device sends a trigger towards a specific UDP port of the device at the other end of the tunnel (i.e. 10.10.10.29). The trigger for VLAN X is sent via the separated control network. If an incoming GVRP registration message for VLAN value X arrives on a non-blocked port of the 10.10.10.23, this indicates completion of a (half) VLAN tunnel. Both these times are logged using the PC clock with a precision of

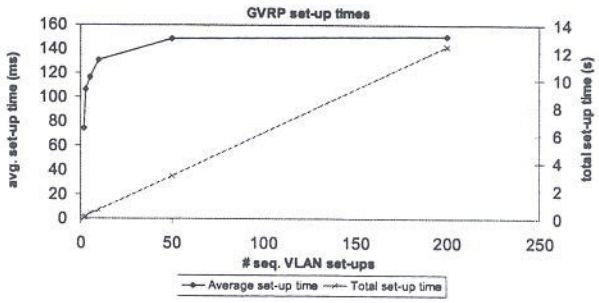

Figure 6: GVRP set-up times (up to 200 VLANs)

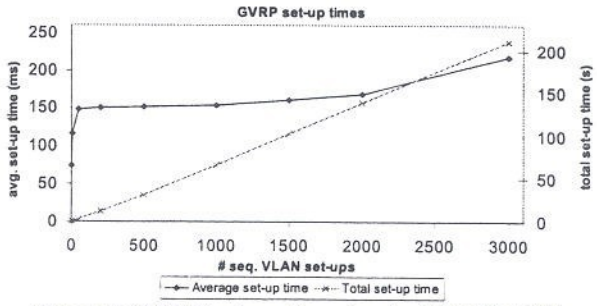

Figure 7: GVRP set-up times (up to 3000 VLANs)

1ms. We measured the set-up times of every VLAN of the burst individually and measured the duration of the entire set-up time. The results (averages of 20 tests) are presented for a growing number of VLANs in Fig. 6. For example, it takes 3028ms to perform 50 sequential VLAN set-ups while the average set-up time measures 148.5ms. This means that the GVRP doesn't handle the requests sequentially. The total set-up time increases linearly with the length of the burst but is not dependent on the individual set-up times. It is dependent on the rate at which the GVRP PDUs are sent. The GVRP inter-arrival time measures nearly a constant rate of 59ms. The GVRP set-up times up to 3000 are represented in Fig. 7. In this range, the set-up times increase slightly faster than linear because the inter-arrival times are not constant in this range but increase from 59ms (50 VLANs) to 70ms (3000 VLANs). Also the average set-up time is no longer constant but starts increasing significantly at bursts of more than 1000 VLANs. For example, it takes 66.5s to set up 1000 VLANs while the average set-up time has increased to 154.25ms.

B. *Scenario 2: Failure Recovery*

If a trunk link fails in the network, all VLANs and associated reservations that are transported on this link, need to be reconfigured. We prefer not to rely on the slow STP reconfiguration and measured the time it takes to establish all these tunnels on a new working path (e.g. on another instance of the STP). The reservation parameters of the failed VLAN tunnels have to be re-invoked on the newly mapped VLAN tunnels. Therefore, the G2RP set-up times are examined. Because the Click Data Plane couldn't cope with heavy configurations consisting of many elements, the Click configuration is kept limited and the associated maximum size of the G2RP bursts is resized accordingly. Fig. 8 represents the G2RP set-up times for a Click configuration of 7 VLANs and 2 QoS classes. If the reserved VLAN 1 is left out, this allows 12

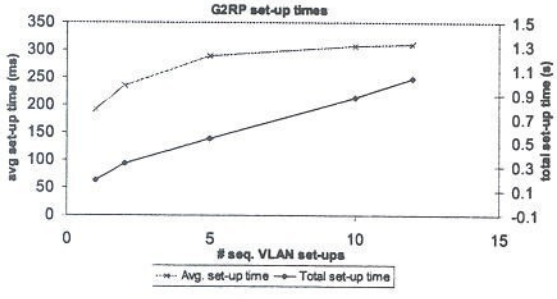

Figure 8: G2RP set-up times

tunnels to be set up. For example, it takes 1033ms for setting up all 12 reservations while the average G2RP set-up time measures 310ms.

C. *Other Performance Measurements*

To correctly evaluate the developed applications, the implementation provides a full-logging option. Because the logging influences the protocol performance, we compared the set-up times of the minimal-logging version and the verbose full-logging version. This is illustrated on Fig. 9. The increasing impact of the logging on the set-up times is clearly noticeable for long bursts of GVRP set-up messages.

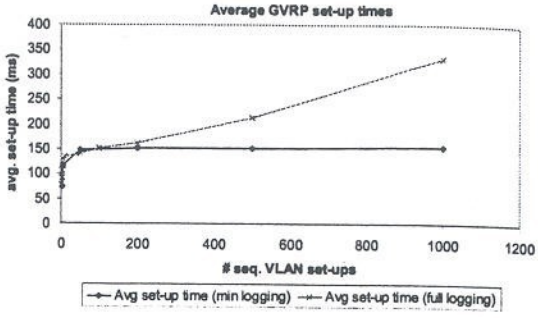

Figure 9: GVRP Set-up time, Min. vs. Full logging

D. *Comparison with other tunnel set-up mechanisms*

As concurrent implementation we take a Linux implementation of RSVP-TE. The set-up times, presented in Table I, are the total set-up times for a single direction of the LSP tunnel in a test network with one intermediate hop. The RSVP-TE times were measured without configuring the switching hardware or without admission control. In these circumstances we can compare these times with the GVRP set-up times of our implementation. Because RSVP is

TABLE I

TOTAL TIMES IN MS TO SET-UP # NUMBERS OF TUNNELS.

Number of tunnels	RSVP-TE[7]	GVRP
10	510	776
50	17350	3849.5
100	63000	7000

suffering from a non-linear growing function in assigning LSP labels, GVRP is faster, even without speed optimizations. However, the Click overhead is clearly noticeable for the lower burst sizes. Compared to these results, the more complex G2RP is also characterized by acceptable (bi-directional) set-up times.

V. CONCLUSION

The presented evaluation strengthens the candidature of Switched Ethernet as a valid access technology and of GVRP/G2RP as a fast and flexible tunnel set-up mechanism.

VI. REFERENCES

[1] Cristèle Bouchat, Sven van den Bosch,Thierry Pollet, Alcatel Bell, *QoS in DSL Access*, IEEE Communications Magazine, p.108-114, Sep 2003.

[2] IEEE 802.1D, Standards for Local and Metropolitan Area Networks: Media Access Control (MAC) Bridges, 1998.

[3] IEEE 802.1q, Standards for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks, 1998.

[4] Meetinghouse Data Communications, *Implementation of GVRP and GARP*, http://www.mgthouse.com, Website.

[5] Fouad A. Tobagi, Pablo Molinero-Fernandez and Mansour J. Karam, *Study of IEEE 802.1p GARP/GMRP Timer Values*, Computer Systems Laboratory, Stanford University, September 1997.

[6] Pim Van Heuven, Frederic Van Quickenborne, Filip De Greve, Brecht Vermeulen, Steven Van den Berghe, Filip De Turck, and Piet Demeester, *VLANs and GVRP on Linux: quickly from specification to prototype using the Click router platform*, Linux Kongress 2003.

[7] Jan Coppens, *Evaluation of recovery in IP/MPLS networks*, Master thesis (in Dutch), June 2001, University of Ghent.

Acknowledgments: Part of this work is supported by the Flemish Government through the IWT project PANEL (IWT020052)

Evaluation of a Tunnel Set-up Mechanism in Qos-aware Ethernet Access Networks

Filip De Greve
INTEC - University of Ghent, Belgium
(filip.degreve@intec.ugent.be)

In collaboration with Alcatel Bell Antwerp, Belgium

1. Generic Attribute Registration Protocol (GARP)

- GARP provides a generic attribute dissemination capability for bridged LANs and the use is described in IEEE 802.1d.
- Two existing implementations are described in IEEE 802.1q:
  - GMRP: GARP Multicast Registration Protocol
  - GVRP: GARP VLAN Registration Protocol

2. Our contribution

- **Our contribution:**

  - For establishment of the tunnels, we designed a "Scoped Refresh" extension of the standardized GVRP
  - For distributing the reservation parameters associated with these tunnels, a new GARP-based protocol is designed. It is called GARP Reservation Parameters Registration Protocol (G2RP).
  - We realized a full implementation of both protocols GVRP and G2RP on a Linux/Click Router test bed .
  - We made extensive performance measurements

3. GVRP Extension

"For establishment of tunnels in a Ethernet-based access network, we developed a "Scoped Refresh" extension of the standardized GVRP. This extension resolves the relevant problem of depletion of the limited sized VLAN Databases."

- "Scoped Refresh" extension

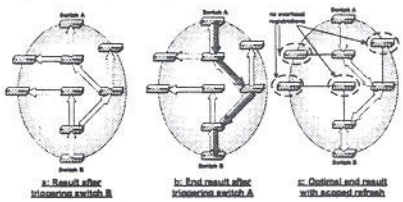

4. Design of new GARP Protocol

"For distributing the reservation parameters associated with the VLAN tunnels, a new GARP application, referred to as G2RP, is developed with its own specific GIP context, closed loop design, mechanism for updating reservation parameters and admission control."

- Closed loop design

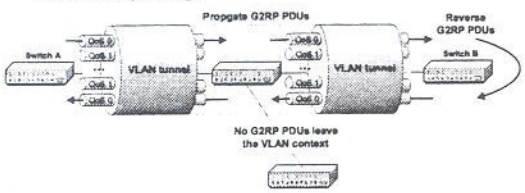

4. Design of new GARP Protocol

- G2RP is designed according to the GARP standard:

  - The GARP broadcast address is assigned: 01-80-C2-00-00-22
  - The G2RP specific attribute type and values are defined.

VLAN-ID	RSV-ID	Reservation Parameters A	Reservation Parameters B

  - After defining the G2RP semantics, the appropriate hardware operations have to be linked with registration and de-registration events.
  - Distributed admission control is added to keep track of the existing bandwidth reservations.
  - Finally, the GIP (= GARP Information Propagation) context is free to define. For the G2RP application, we linked the GIP context to the pre-established VLAN tunnel

5. Implementation Aspects

- Both protocols, GVRP and G2RP, are implemented and tested on a Linux test bed.

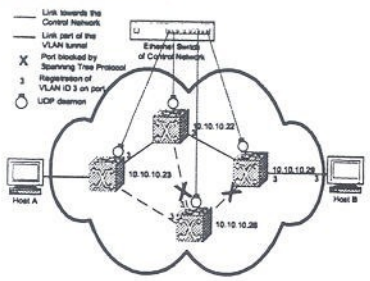

5. Implementation Aspects

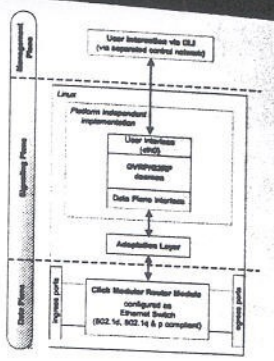

6. Evaluation results

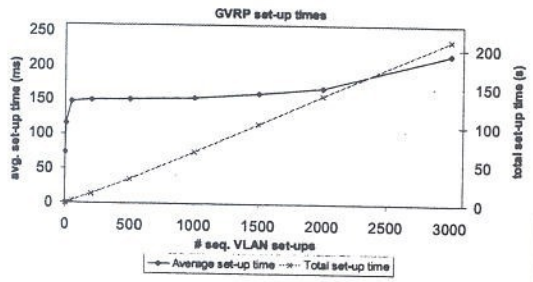

6. Evaluation results

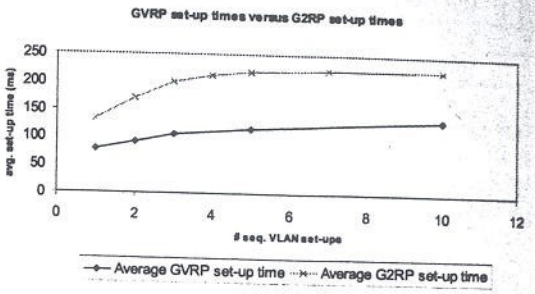

6. Evaluation results

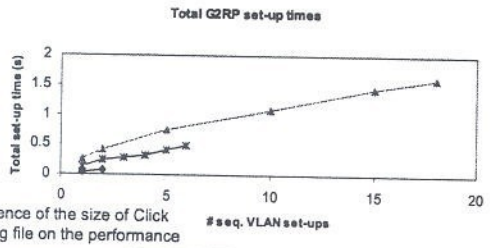

Influence of the size of Click config file on the performance

438 lines ← 7 VLAN/2QoS/6 — 4VLAN/2QoS → 294 lines
582 lines ← 10VLAN/2QoS/2 — 2VLAN/2QoS → 198 lines

7. Comparison with RSVP

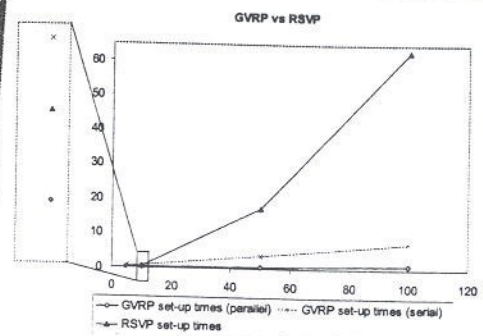

LANMAN 2004

The 13th IEEE Workshop on Local and Metropolitan Area Networks

The Acqua Hotel

555 Redwood Hwy, Mill Valley, CA
(San Francisco Bay Area)
25 - 28 April, 2004

Sponsored by

IEEE

IEEE COMMUNICATIONS SOCIETY

Patrons

hp invent

NOKIA
CONNECTING PEOPLE

13th IEEE Workshop on Local and Metropolitan Area Networks

Wednesday, April 28, 2004

Evaluation of a tunnel set-up mechanism in QoS-aware Ethernet access networks

De Greve, F. Van Quickenborne, F. Van Heuven, P. De Turck, F. Vermeulen, B. Van den Berghe, S. Moerman, I. Demeester, P. Van den Bosch, S. Janssens, N. Bouchat, C. Van Vlerken, B.
Dept. of Inf. Technol., Ghent Univ., Belgium;

This paper appears in: **Local and Metropolitan Area Networks, 2004. LANMAN 2004. The 13th IEEE Workshop on**
Publication Date: 25-28 April 2004
On page(s): 247- 252
ISSN:
ISBN: 0-7803-8551-9
INSPEC Accession Number: 8169270
Digital Object Identifier: 10.1109/LANMAN.2004.1338440
Posted online: 2004-10-08 13:18:24.0

Abstract
This paper describes the design of a novel tunnel set-up mechanism for Ethernet access networks. The establishment of VLAN-based tunnels is realised with a "scoped refresh" extension of the GVRP standard. For the distribution of QoS-related reservation parameters, a new GARP-based protocol (G2RP) is developed with a closed loop design for updating the reservation parameters and hop-by-hop admission control. Both protocols are implemented on a Linux test bed using the click modular router. Extensive performance evaluations were performed. The results are compared with Linux implementations of the other tunnel set-up mechanisms.

Index Terms
Inspec

Controlled Indexing
access protocols local area networks quality of service subscriber loops telecommunication congestion control telecommunication network routing

Non-controlled Indexing
QoS-aware Ethernet access networks click modular router hop-by-hop admission control protocols quality of service reservation parameters tunnel set-up mechanism virtual local area networks

Author Keywords
Not Available

References

No references available on IEEE Xplore.

Citing Documents

No citing documents available on IEEE Xplore.