

VRRP Working Group
INTERNET DRAFT
[draft-ietf-vrrp-lane-02.txt](#)

Atul Bansal
(FORE Systems)
Rob Enns
(Juniper Networks)
Vivek Menezes
(Pluris)
Rob Montgomery
(Cabletron Systems)
Ayikudy K.Srikanth
Hamayon Mujeeb
(Nortel Networks)
Expires June 2000

VRRP Operation over ATM LAN Emulation

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-Draft can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

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

Abstract

This memo specifies the application of the Virtual Router Redundancy Protocol (VRRP) over networks built using the ATM Forum LAN Emulation protocol (LANE). The memo specifies the behavior required of a LANE client (LEC) in order to interoperate with VRRP.

Two schemes "Proxy LEC Scheme" and "Non-Proxy LEC Scheme" are outlined for running VRRP over LANE. Both schemes can be deployed

simultaneously in the network.

1. Introduction

The Virtual Router Redundancy Protocol (VRRP) [[1](#)] specifies a protocol used to provide dynamic failover in forwarding responsibility for a set of IP addresses. VRRP provides high availability of a default forwarding path without requiring hosts to run routing or router discovery protocols.

ATM Forum LAN Emulation (LANE) [[2,3](#)] specifies a protocol used to emulate Ethernet/IEEE 802.3 and Token Ring/IEEE 802.5 LANs on an Asynchronous Transfer Mode (ATM) network. LANE is used to extend LAN connectivity over an ATM backbone.

This memo specifies the behavior required of a LANEv2 client in order to interoperate with VRRP. The behavioral difference between the LANEv2 client and LANEv1 client is described in [section 7](#) so that a LANEv1 client can interoperate with VRRP.

2. Discussion

VRRP operates by electing a master router which responds to a well known virtual MAC address (VMAC) for the virtual router. When a new master router is elected, the physical device responding to the VMAC may change. LANE operates by maintaining a set of mappings from MAC addresses to ATM addresses. When running VRRP on a LANE network, the LANE protocol must be notified when a new master router is elected so that it can update the MAC address to ATM address mapping within the ELAN for the virtual router's MAC address. In essence while running VRRP over LANE a virtual MAC address may change location from one LEC to another.

Some Definitions:

1) ADDRESS ASSOCIATION:

When a LEC represents a particular virtual MAC address the virtual MAC address is "associated" with the LEC. No two LECs should be associated with the same MAC address in an ELAN.

2) ADDRESS DISASSOCIATION:

When a LEC stops representing a virtual MAC address which it apriori represented it is said to have disassociated itself from the virtual MAC address.

3) Proxy LEC

A LEC which doesn't register its MAC address associations with the LES. A proxy LEC is responsible for replying to LE_ARP requests for the MAC addresses it is associated with.

4) Non-Proxy LEC

A LEC that registers its MAC address associations with the LES. The LES is responsible for replying to LE_ARP requests for MAC addresses registered by a non-proxy LEC.

5) Targetless LE_ARP_REQUEST:

As defined in section 7.1.30 of [3].

"An LE Client MAY generate an LE_ARP_REQUEST with the TARGET-LAN-DESTINATION field indicating not present, but with the SOURCE-LAN-DESTINATION, SOURCE-ATM-ADDRESS, and optionally, the LLC-Muxed-ATM-Address TLV present. This allows the generating LE Client to advertise an LE_ARP cache binding to all other LE Clients in its ELAN. An LE Client receiving a targetless LE_ARP_REQUEST MUST delete any LE_ARP cache bindings to that SOURCE-LAN-DESTINATION, and MAY replace any such bindings with the information from the LE_ARP_REQUEST.

If an LE Client sends out a targetless LE_ARP_REQUEST for a currently registered unicast LAN Destination, it must have the same bindings and TLVs as the most recent LE_REGISTER_REQUEST for that LAN Destination sent to the LE Server.

An LE Client receiving a targetless LE_ARP_REQUEST for a SOURCE-LAN-DESTINATION not in its cache SHOULD ignore the request. An LE Client receiving a targetless LE_ARP_REQUEST for a SOURCE-LAN-DESTINATION that is in its cache SHOULD update that cache entry with the information in the LE_ARP_REQUEST. This technique ensures that the LE Client has up-to-date information, without filling its LE_ARP cache entries for which it has no use."

6) No-source LE_NARP_REQUEST:

As defined in section 7.1.31 of [3].

"An LE Client MAY generate a no-source LE_NARP_REQUEST to notify other clients that it no longer represents the unicast LAN Destination in TARGET-LAN-DESTINATION and TARGET-ATM-ADDRESS (and LLC-Muxed-ATM-Address TLV, if present). In this case, SOURCE-ATM-ADDRESS MUST be zero." On receiving a No-Source LE_NARP_REQUEST a LEC should remove the TARGET-LAN-DESTINATION to TARGET-ATM-ADDRESS mapping from its ARP cache.

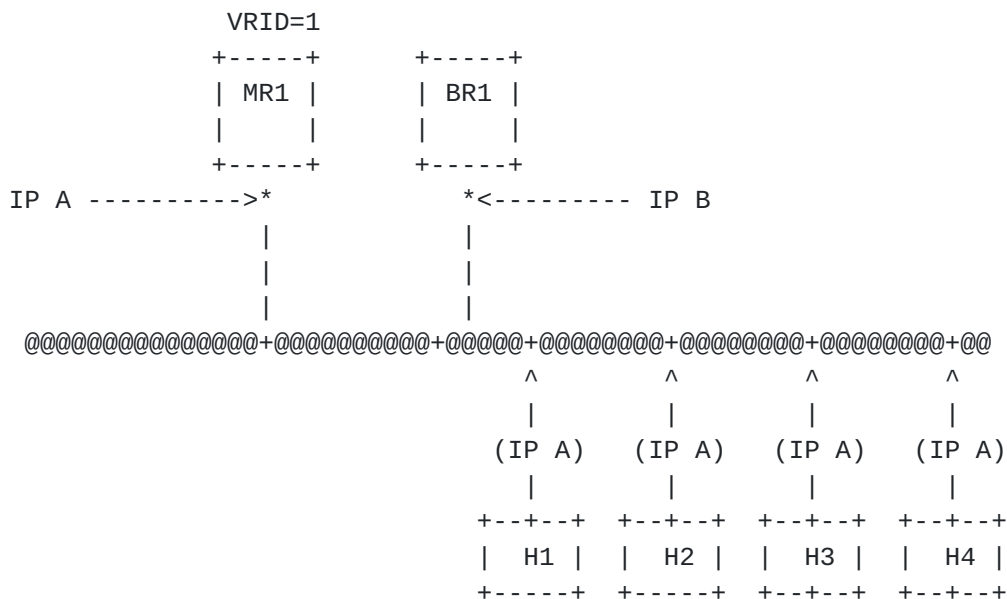
In VRRP virtual MAC addresses can change their associations with LEC.

Since there are two different LEC configurations in an ELAN (Proxy and Non-Proxy) we discuss ADDRESS ASSOCIATION with both types.

Two common configurations are envisioned when running VRRP over LANE. In the first, the LANE client (LEC) is co-located with the router running VRRP. In other words, the LANE client represents one of the router's network interfaces. In the second scenario, a router attached to a LAN such as Ethernet is located behind a bridge running a LANE client. The scenarios are considered in detail in the following sections.

2.1 Scenario 1 ATM only Attached VRRP Routers

In this scenario we assume the all routers running VRRP are attached to the LANE network. This scenario is analogous to the first sample configuration in [1], replacing the Ethernet/Token Ring/FDDI network with a LANE network.



Legend:

```

@@@+@@@+@@@+@@ = LANE network
H = Host computer
MR1 = Master Router
BR1 = Backup Router
* = IP Address
(IP) = default router for hosts

```

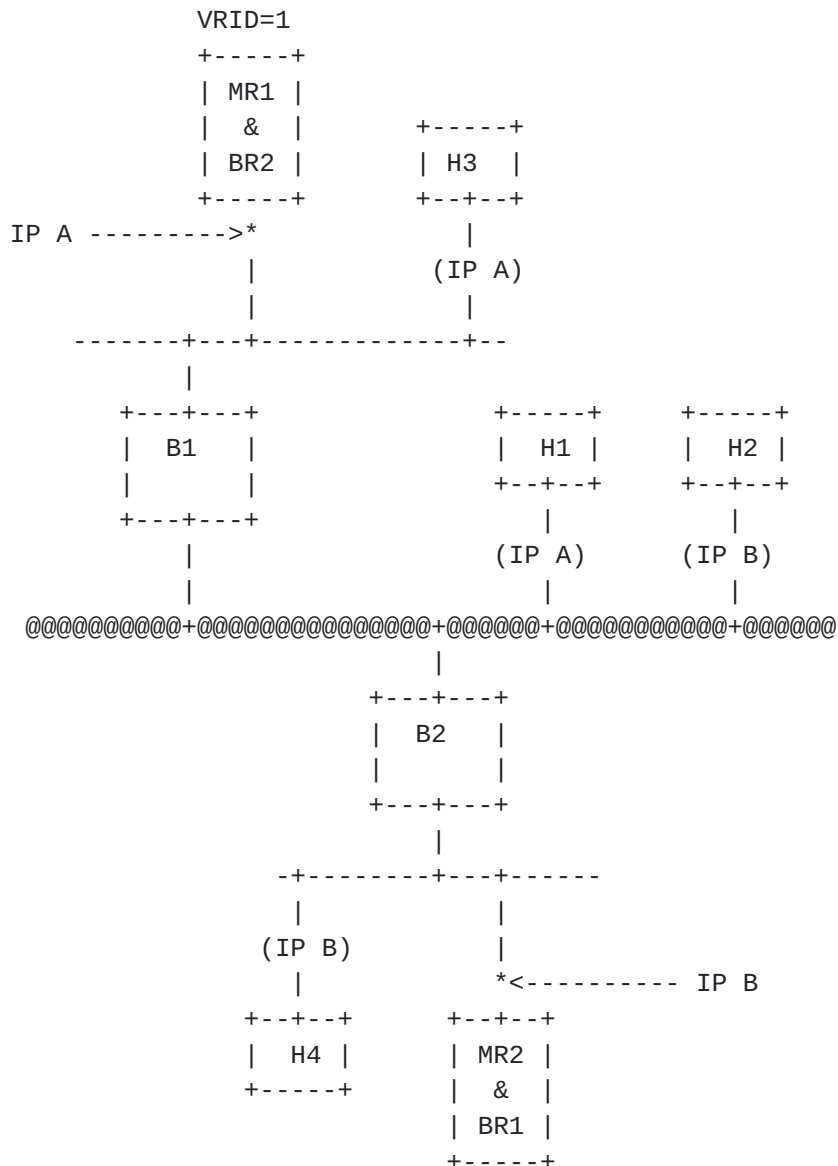
In the above configuration, the end-hosts install a default route to virtual router #1's IP Address (IP A). All the end-hosts (H1-H4) bind (IP A) with the virtual router's (VRID=1) virtual MAC address

(VMAC1) which in turn is bound to MR1's link layer address (MR1 LEC's ATM Address).

Upon MR1 failure, BR1 becomes the master and sends out VRRP advertisement and its LEC ASSOCIATES itself with VMAC1.

2.2 Scenario 2 ATM and LAN Network Attached

In this scenario we assume that all routers running VRRP are attached to the LAN network behind the bridges. Furthermore, there are two virtual routers backing each other. Half of the hosts are using one virtual router and the other half of the hosts are using the second virtual router.



VRID=2

Legend:

```

@@@+@@@+@@@+@@ = LANE network
---+---+---+--- = Ethernet network
      H = Host computer
MR1, MR2 = Master Router
BR1, BR2 = Backup Router
B1,B2 = LANE Bridges
      * = IP Address
(IP) = default router for hosts

```

In the above configuration, MR1 is the master for virtual router #1 serving IP address (IP A) and MR2 is the master for virtual router #2 serving IP address (IP B). The end-hosts H1 and H3 install a default route to the IP address of virtual router #1 (IP A) and bind (IP A) with VMAC of the virtual router #1 (VMAC1) which in turn is bound to the Bridge1's ATM Address (VMAC1 is behind bridge 1). Similarly, end-hosts H2 and H4 install a default route to (IP B) and bind VMAC2 to Bridge2's ATM address.

Upon MR1 failure, BR1 becomes the master and sends out the VRRP advertisement with VMAC1 as the source MAC address. The VRRP advertisement triggers Bridge 1 and Bridge2 to update their station learning cache. Upon detecting the station move, bridge B1 makes an ADDRESS DISASSOCIATION from VMAC1 and bridge B2 makes an ADDRESS ASSOCIATION with VMAC1

Another topology in Scenario 2 can have ATM attached as well as LAN attached VRRP Routers. In this case, the ATM attached VRRP Router can use either Proxy or Non-Proxy LEC Scheme and the LANE bridges use Proxy LEC Scheme.

3.0 Detailed Description of VRRP over LANE

3.1 Non-Proxy LEC Scheme

This scheme is applicable when an ADDRESS ASSOCIATION or ADDRESS DISASSOCIATION is made by a Non-Proxy LEC.

A LEC joins the ELAN as a non-proxy client. It registers with LES the physical MAC address associated with the ELAN interface.

On having to make an ADDRESS ASSOCIATION with VMAC the LEC registers the VMAC with the LES and sends out a "Targetless LE_ARP_REQUEST"

with the VMAC to its own ATM address binding. If the VMAC registration fails, the master must persist with the registration phase. The LEC continues to send the Targetless LE_ARP_REQUESTS (subject to the generation time limits defined in [3]) every second. (In case of LANE v1 client, send LE_NARP with VMAC as the TARGET-LAN-DESTINATION and with zero TARGET-ATM-ADDRESS and zero SOURCE-ATM-ADDRESS)

If a Non-Proxy LEC makes an ADDRESS DISASSOCIATION it unregisters the VMAC and sends out "No-Source LE_NARP" with the VMAC to its own ATM address mapping.

If the VRRP router fails without being able to make an ADDRESS DISASSOCIATION, the LES detects the Control VC failure to the LEC associated with the VMAC and deletes the current VMAC to ATM binding. Other LECs in the ELAN will lose their VCs to the failed VRRP router. There is a potential failure mode if the control VC to the LES isn't dropped when the master router fails.

Note that the ATM switch can only detect User-Network Interface (UNI) failure and not the LEC failure. The ATM UNI signalling protocol [4] specifies a mechanism for detecting the interface failure between the ATM switch and the ATM host. This mechanism relies on receiving the AAL failure notification from the Signalling ATM Adaption Layer (SAAL). Upon receiving the AAL failure notification, UNI signalling waits for the duration specified by the T309 timer value (default is 4 secs) before releasing all active calls on the failed interface.

The SAAL layer [5,6,7] provides reliable transport of signalling messages between peer entities (e.g., ATM switch and host) and provides mechanism for detecting the failure between the peer entities. Upon detecting the failure between the peer entities, the SAAL layer notifies the signalling Layer via AAL failure notification.

Some ATM switch implementations release all VCs on the failed interface upon link carrier loss. This memo recommends to use release on carrier loss feature if available. Otherwise, this memo recommends that both UNI and SAAL timers must be set to detect the UNI failure under 3 seconds. Upon the UNI failure detection, the ATM switch must immediately release all VCs related to the failed link.

3.2 Proxy LEC Scheme

This scheme is applicable to a Proxy LEC that makes ADDRESS ASSOCIATION and ADDRESS DISASSOCIATION.

The LEC must register as a proxy client and can optionally register its own MAC address with a LES (for pings, TELNET, SNMP, etc.).

On making an ADDRESS ASSOCIATION a LEC sends out a "Targetless LE_ARP_REQUEST" with the VMAC to its own ATM address binding. It also makes LE_ARP_REQUEST for this VMAC address. If it gets a successful reply then it should persist sending "Targetless LE_ARP_REQUEST" until no reply is received. This retransmission should not occur more frequently than once per second. (In case of LANE v1 client, send LE_NARP with VMAC as the TARGET-LAN-DESTINATION and with zero TARGET-ATM-ADDRESS and zero SOURCE-ATM-ADDRESS)

When the LEC makes an ADDRESS DISASSOCIATION it sends out "No-Source LE_NARP" with the VMAC to its own ATM address mapping.

3.3 State Description

This section augments the state description defined in Sec 6.4 of [1].

3.3.1 Initialize

Register either as a Proxy LEC or as a Non-Proxy LEC and complete the LANE JOIN PHASE [3].

- If the Priority = 255
 - o Make an ADDRESS ASSOCIATION with the VMAC
 - o Send an ADVERTISEMENT
 - o ...(same as [1])
 - o Transition to the {Master} State
- else
 - o ...(same as [1])
 - o Transition to the {Backup} State

3.3.2 Backup

While in this state, a VRRP router must do the following:

- MUST not respond to ARP requests for the IP address(es) associated with the virtual router.
- MUST not respond to LE_ARP_REQUEST for VMAC
- ...(same as [1])
- If the Master_Down_Timer fires, then
 - o Make an ADDRESS ASSOCIATION with the VMAC.


```
    o send an ADVERTISEMENT
    o Broadcast a gratuitous ARP .....
    o ...(same as [1])
    o Transition to the {Master State}
endif
```

3.3.3 Master

While in the {Master} State the router functions as the forwarding router for the IP address(es) associated with the virtual router.

While in this state, a VRRP router must do the following:

- MUST respond to ARP requests for the IP address(es) associated with the virtual router
- MUST respond to LE_ARP_REQUEST for VMAC
- ... (same as [1])
- If a shutdown event is received, then
 - o cancel the Adver_Timer
 - o send an ADVERTISEMENT with Priority = 0
 - o Make an ADDRESS DISASSOCIATION with the VMAC.
 - o Transition to {initialize} state
- endif
- ... (same as [1])
- If an ADVERTISEMENT is received, then
 - ... (same as [1])
 - If the Priority in the ADVERTISEMENT is greater than the local Priority
 - or
 - If the Priority in the ADVERTISEMENT is equal to the local Priority and the the Primary IP address of the sender is greater than the local primary address, then:
 - o Make an ADDRESS DISASSOCIATION with the VMAC.
 - o Transition to {Backup} state

4.0 Requirement on the LANE Bridges for the VRRP Operation over LANE

A LANE bridge must register as a proxy client with LES. This is the default for LANE bridges. A VRRP router could attach on a LAN port of a LANE Bridge.

A LANE bridge must make ADDRESS ASSOCIATION on an ELAN port with MAC addresses that it learns on other ports. It must make ADDRESS DISASSOCIATION when it learns an a MAC address on an ELAN port on which it has made an ADDRESS ASSOCIATION with the same MAC.

The ADDRESS ASSOCIATION and ADDRESS ASSOCIATION by LANE bridges upon detecting the station move is essential for the correct operation of VRRP in various extended LAN topologies. This behavior should be default behavior for all LANE bridges independent of VRRP.

5.0 VRRP operation over Token Ring LANE networks

The operation of VRRP over LANE Token Ring is similar to the Ethernet LANE case discussed above. However, as discussed in [1] [RFC 2338](#), VRRP over token ring requires the use of a token ring functional address for the VRID virtual MAC address because of the absence of a general multicast mechanism over old and new token ring adapter implementations. The token ring functional address support is the only generally available multicast mechanism.

The hosts on a Token Ring LANE subnet will use the Token Ring functional address as the destination MAC address for all IP traffic going through the IP address(s) of the virtual router. As the functional address is a multicast address, all such traffic goes through the ATM LANE BUS, causing a potential congestion of the BUS.

VRRP implementations over pure Token Ring LANE networks (Scenario 1) should use IEEE 802 MAC Address used for Ethernet, that is, 00-00-5E-00-01-{VRID}. In this case, the lack of a real token ring adapter makes the need for a functional address a non issue. In the case of token ring bridged environment, as long as the Master and all Backups are on LANE, hosts can be on LANE as well as on Token Ring and there is no need to use functional addresses.

However, if the Master or any of the Backups is on a bridged Token Ring in a mixed network of bridged token ring and token ring emulation (Scenario 2), the use of functional addresses cannot be avoided.

In the case of VRRP in a Token Ring Source route bridged LANE environment, either the Virtual MAC address or the functional address could be used but the route descriptor should not be configured.

6.0 Behavior difference between the LANE V1 vs. LANE V2 clients

From 7.1.29 [3]. "An LE Client MAY generate an unsolicited LE_NARP_REQUEST to allow other clients to learn about a changed Remote LAN-ATM address binding. This LE_NARP_REQUEST advertises that the generating LE Client believes that an old binding between TARGET-LAN-DESTINATION and TARGET-ATM-ADDRESS is no longer valid, and that the generating LE Client now serves the LAN destination. An LE

Client receiving a valid LE_NARP_REQUEST MAY delete any of its LE_ARP cache entries that match the TARGET-LAN-DESTINATION and TARGET-ATM-ADDRESS, and MAY replace such entries with the binding between TARGET-LAN-DESTINATION and SOURCE-ATM-ADDRESS. The generating LE Client MUST include the ATM address of the LE Client which was previously representing the LAN destination. The LLC-Muxed ATM Address TLV MUST NOT be included in a LANE v1 LE_NARP_REQUEST."

The V1 LE_NARP_REQUEST message didn't consider the two separate cases of updating ATM address to MAC binding by LANE clients. The first case is where a LEC who was serving the MAC address by providing the MAC to its own ATM address. Now, the LEC does not representing the MAC and it has to inform other LECs about this fact. The second case is where a LEC starts serving the MAC address and has to inform other LECs about this fact.

LANEv2 defines two separate messages to handle these two cases. "No-Source LE_NARP_REQUEST" handles the first case and "Targetless LE_ARP_REQUESTS" handle the second case.

LANEv1 LE_NARP_REQUEST as defined in [1] can only be used by the client for the second case. The VRRP operation with LANEv1 requires that the LANEv1 client send the LE_NARP_REQUEST with zero TARGET-ATM-ADDRESS and zero SOURCE-ATM-ADDRESS. This is change from [2] and [3]. The LANEv1 and LANEv2 LES must forward this LE_NARP_REQUESTS to all LECs. Upon receiving LE_NARP with the Target ATM address of 0x00, all clients (proxy and non-proxy) should delete the MAC to ATM address mapping for the MAC listed in the TARGET-LAN-ADDRESS.

7.0 Security Consideration

This memo does not define any new protocol packets. Therefore, it follows the same security mechanisms as defined in [1,2,3].

8.0 Acknowledgments

The authors would like to thank Nischal Sheth, Mike Kazar, Yevgeny Shakhnovich, Geetha Brown and Danny Mitzel for their comments and suggestions.

References

- [1] Virtual Router Redundancy Protocol, [RFC 2338](http://ftp.ietf.org/rfc/rfc2338.txt), Internet Engineering Task Force, April 1998.
[ftp://ftp.ietf.org/rfc/rfc2338.txt](http://ftp.ietf.org/rfc/rfc2338.txt).

- [2] LAN Emulation over ATM Version 1.0, The ATM Forum, January, 1995.
<ftp://ftp.atmforum.com/pub/approved-specs/af-lane-0021.000.pdf>.
- [3] LAN Emulation over ATM Version 2.0 - LUNI Specification, The ATM Forum, July, 1997.
<ftp://ftp.atmforum.com/pub/approved-specs/af-lane-0084.000.pdf>.
- [4] The ATM Forum, ATM User-Network Interface (UNI) Signalling Specification, Version 4.0, July 1996.
<ftp://ftp.atmforum.com/pub/approved-specs/af-sig-0061.000.pdf>
- [5] ITU-T Q.2391 (Note 1), B-ISDN User-Network Interface Layer 3 for Basic Call/Control, March 1994
- [6] ITU-T Q.2110, B-ISDN Signalling ATM Adaption Layer - Service Specific Connection Oriented Protocol (SSCOP) (Note)
- [7] ITU-T Q.2130, B-ISDN Signalling ATM Adaption Layer - Service Specific Coordination Function for support of signalling at the user-network interface (SSCF at UNI) (Note)

Authors' Addresses

Atul Bansal
8221 Bell Lane
Vienna, VA 22182 USA
Phone: +1 703 876 3810
Email: atulbans@aol.com

Vivek Menezes
Pluris
10455 Bandley Dr,
Cupertino, CA 95014 USA
Phone: +1 408 861 4147
Email: vivek@pluris.com

Rob Enns
Juniper Networks
385 Ravendale Drive
Mountain View, CA 94043 USA
Phone: +1 650 526 8000
Email: rpe@juniper.net

Rob Montgomery
Cabletron Systems
Pease Training Center
Box 5005, 35 Industrial Way
Rochester, NH 03866-5005 USA

INTERNET DRAFT

Expires June 2000

Phone: +1 603 337 9258
EMail: rmontgom@cabletron.com

A.K.Srikanth
Nortel Networks Inc
600 Tech. Park Drwy.
Billerica, MA 01821 USA
Phone: +1 978 916 8226
EMail: asrikant@baynetworks.com

Hamayon Mujeeb
Nortel Networks Inc
600 Tech. Park Drwy.
Billerica, MA 01821 USA
Email: hmujeeb@nortelnetworks.com

