

Internet Draft
Document: [draft-ietf-forces-discovery-00.txt](#)
Expires: April 24, 2005
Working Group: ForCES

Furquan Ansari
Lucent Tech.
Hormuzd Khosravi
Intel Corp.
Jamal Hadi Salim
Znyx
October 25, 2004

ForCES Intra-NE Topology Discovery

[draft-ietf-forces-discovery-00.txt](#)

Status of this Memo

By submitting this Internet-Draft, I certify that any applicable patent or other IPR claims of which I am aware have been disclosed, and any of which I become aware will be disclosed, in accordance with [RFC 3668](#).

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/lid-abstracts.txt>.

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

This Internet-Draft will expire in April 24, 2005.

Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC-2119](#)].

Abstract

This document describes a mechanism for discovering inter-FE topology and topology maintenance. Such a mechanism is essential for
Ansari et al. Expires: Jan 2005 [Page 1]

all these elements in the set to behave as a single Network Element, as required by the ForCES architecture as well as to perform certain optimizations at the FE by making use of the topology. The discovery mechanism only operates during post-association phase of ForCES protocol.

Table of Contents

- [1. Definitions.....2](#)
- [2. Introduction.....2](#)
 - [2.1. Motivation.....4](#)
- [3. Topology Discovery Mechanism.....5](#)
 - [3.1. Minimum requirements.....6](#)
 - [3.2. Protocol Details.....6](#)
 - [3.2.1. Topology Discovery and Maintenance.....7](#)
 - [3.2.2. Full topology computation at the CE from partial topologies.....8](#)
 - [3.3. Protocol and Message Headers.....98](#)
 - [3.3.1. TLV definitions.....10](#)
 - [3.4. Inter-FE Topology Discovery Examples.....10](#)
 - [3.4.1. Forwarding Elements connected in a daisy chain.....11](#)
 - [3.4.2. Forwarding Elements connected in a ring.....12](#)
- [4. Security Considerations.....13](#)
- [5. References.....13](#)
 - [5.1. Normative.....13](#)
 - [5.2. Informative.....14](#)
- [6. Authors' Addresses.....14](#)
- [7. Full Copyright Notice.....14](#)
- [8. Acknowledgements.....15](#)

[1. Definitions](#)

Inter-FE topology discovery: Topology discovery relates to how the FEs are interconnected with each other with respect to packet forwarding. This is the complete view of the intra-NE network as seen by the CE.

Inter-FE topology maintenance: Once the inter-FE topology has been discovered, it has to be continuously monitored to ensure that any changes to the topology are reported to the corresponding CE. This represents the steady state and final phase of the protocol.

[2. Introduction](#)

The ForCES framework document [[RFC 3746](#)] describes how a set of control elements (CEs) and forwarding elements (FEs) interact with each other to form a single network element (NE). It describes the ForCES post-association phase protocol working across the Fp reference point between CE and FE. This document describes an important aspect of the ForCES operational infrastructure, that of discovering the layout of the different elements within an NE.

We describe a mechanism for obtaining the Intra-NE/Inter-FE topology.

The mechanism is divided into two distinct operational pieces:

- . The FE side component that collects FE neighbor information.
- . The CE side component that uses the neighbor information to compute the NE topology.

Given the above split, we believe that this mechanism fits well within a description of Topology Discovery LFB.

The mechanism at the FE can be divided into two modes or phases.

- . Phase I corresponds to the actual discovery process wherein each element discovers its neighbor and maintains a neighbor relationship. Upon Joining an NE, the CE will instruct the FE to start collecting this information. This happens when the FE is administratively up and the associated neighbor links are deemed to be provisioned and operationally up.
- . Phase II corresponds to the topology maintenance phase, wherein any changes to the inter-FE topology during normal operation is reported to the corresponding CE so that an updated view is available. The CE then makes all services it is controlling aware of such details. As an example, based on policy, in a basic IPV4 service, the tables of associated FEs will need to be updated.

As noted above, both phases occur during the post-association phases of the ForCES protocol. In other words, the ForCES protocol association between the CE and the FEs should already have taken place for the discovery mechanism to kick in. This is required because the neighbor relationships maintained by the FEs are reported back to the CE (or queried by the CE) over the ForCES protocol.

The proposed discovery mechanism is required to scale to a very large number of forwarding elements in the NE, with minimal impact on the resources. The following list provides some of the features and goals of the discovery mechanism.


- . Determine connectivity between elements
- . React to changes in link connectivity

- . Construct topology information from the collected partial topology information
- . Tolerant to protocol message losses
- . Applicable to all inter-FE network topologies such as ring, mesh, star etc.
- . Cause minimal overhead
- . Agnostic of the network interconnect technology

2.1. Motivation

The ForCES architecture defines a network element (NE) as a single managed entity made up of a collection of FEs and CEs and is indistinguishable from other network elements in the network. This NE model definition leads to three types of links from the network's perspective: internal (or intra-NE) links and external (or inter-NE) links and control links. Intra-NE links are purely internal to the NE and are not exposed to the external world; whereas, inter-NE (or external) links are exposed to the external world and over which routing adjacencies (such as OSPF, IS-IS, BGP etc.) can be formed. An NE can contain FEs that have zero or more internal/external links ; e.g. in Fig. 1, FE3 has two internal links and no external links while FE1 and FE2 have two internal links and one external link each. Control links are those links that are used for communication between the CE and FE. If the CE and FE are a single Layer 3 hop away as in Fig.1, the control link is typically a physical link e.g. link A of FE1 in the figure. Control links can be logical as well.

A packet entering a ForCES NE may travel multiple FEs within the NE before it exits onto the output link. This requires that the packet be correctly forwarded from the ingress FE to the egress FE. This internal forwarding requires knowledge of the physical FE inter-connection topology so that the CE can appropriately setup internal LFB tables at each FE to handle packet traversal in a sane manner. We use a simple topology discovery mechanism that only operates on internal links and provides the necessary routing information to forward the packets from the ingress FE to the egress FE. Further, the mechanism should be able to reroute around internal link failures, if a path exists. This makes the NE highly available and resilient.



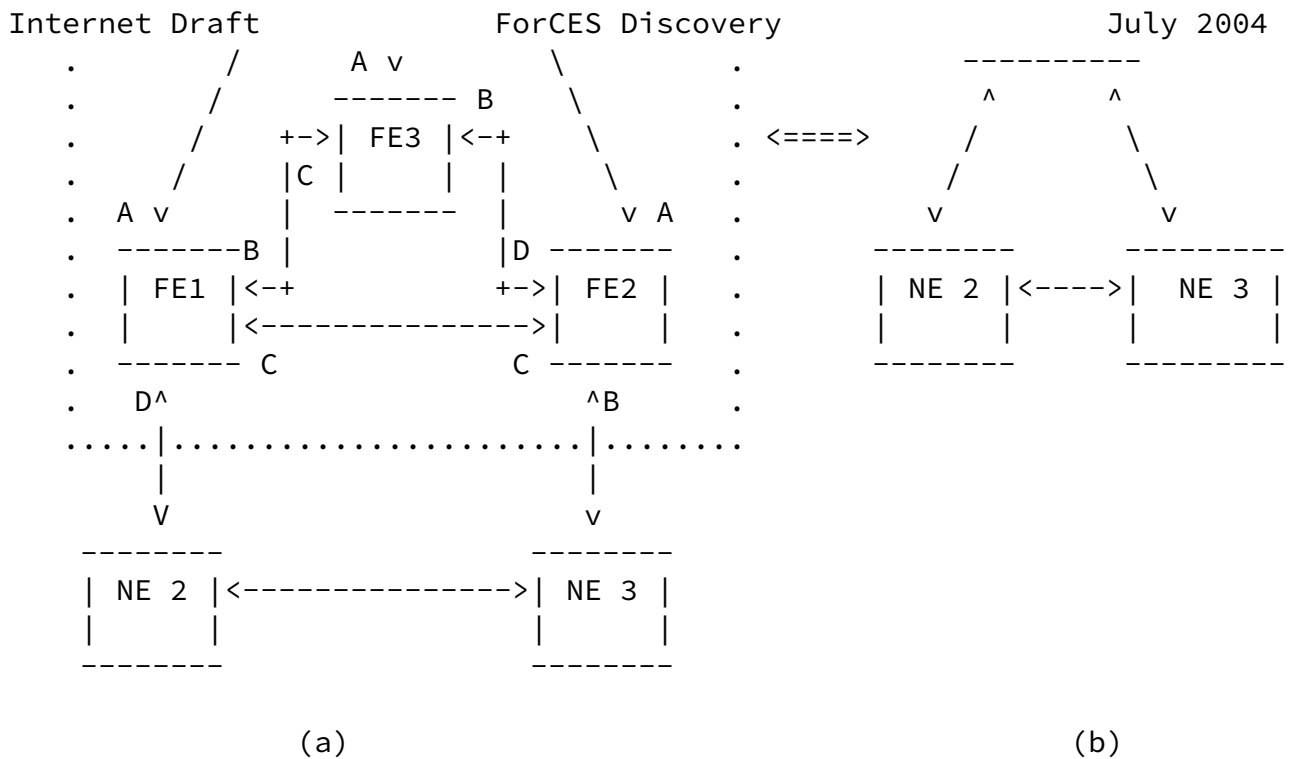


Figure 1:(a) illustrates the internal/external links and topology within a NE. (b) Shows the network topology as seen by external routing protocols

3. Topology Discovery Mechanism

The topology discovery mechanism described here will be restricted to the case where the control and the forwarding elements are a single layer 3 hop away. However, there is no restriction on the number of layer 2 hops between the CE and the FEs. Although, the mechanism can be extended to the multi-hop scenario, it is considered beyond the scope of this document to describe it. The mechanism is expected to work on point-to-point as well as multi-access links.

In order to keep the discovery and maintenance mechanism as simple as possible, the FEs only maintain relationships with the respective neighbors to determine the status of the neighbors. No databases are exchanged between the neighbors. This implies that the topology view for each FE is only limited to the adjacent elements. This partial topology information is reported back to the CE (or queried by the CE) over the ForCES protocol. Since the CE receives such information from all the FEs, it can easily construct the full topology from individual partial topologies reported by each FE. Once the CE

constructs the full topology, such information can be passed to the FEs, if needed (depending on policy).

Internet Draft

ForCES Discovery

July 2004

Topology information is needed by a lot of LFBs and associated services that span multiple FEs within a NE. In the case where the FE aids the CE in offloading the table updates, then it makes sense for the FE to be topology aware. It is sometimes also helpful to keep full topology information at the FEs for cases such as message snooping optimizations. For example, if an FE is aware of the topology, it could snoop on messages sent to other FEs (e.g. broadcasts, multicasts) and update its own tables dynamically without involving the CE. Another example would be FE-FE primary-backup handover scenario. With each FE being fully aware of the complete topology, the backup FE can take over the responsibilities of the primary without involving the CE for such a handover.

[3.1. Minimum requirements](#)

In order for the protocol to work as described, the following assumptions are made.

- . Each element has been configured with their respective IDs (CEID, FEID)
- . Element bindings process has already taken place. In other words, the CE know all the FEs it wants to control and each FE knows which CE is allowed to control it.
- . The ForCES protocol association has already taken place between the CE and the FE in question.
- . The protocol is enabled on the required interfaces.

Note that these are configuration requirements and are satisfied by the respective managers (CEM/FEM).

[3.2. Protocol Details](#)

Once the ForCES protocol association has been established between a CE and a given FE ; i.e. it is in post-association phase, the CE starts sending/advertising Hello/Probe messages to the FE's neighbors such that the messages go through the given FE. In other words, it looks like the given FE is generating probe messages to the neighbor (except that these messages are coming from the CE over the ForCES protocol first). However, this functionality of generating probe messages by the CE can be offloaded to the FE itself (to be more precise, to an FE LFB) ; so that the FE can originate and terminate the probe messages. This provides better scalability of the CE and its resources. The CE can now simply query each FE's neighbor relationship database and register for any

events related to topology changes.

All Hello/Probe messages travel a single PE hop and are not routed to other elements beyond the first hop. This is ensured by using a TTL of one on all Hello/Probe packets. The messages are sent as IP datagrams (multicast/broadcast, where applicable, or unicast in Ansari et al. Expires: Jan 2005 [Page 6])

Internet Draft ForCES Discovery July 2004

general) to the neighboring elements over configured interfaces. Each FE topology LFB component maintains the neighbor relationships as long as the Hello messages are received from the neighbor. If it does not receive a pre-determined number (configured) of back-to-back Hello messages from a given neighbor, it deletes the entry from the database and reports this change to the CE in the form of an event-driven message over the ForCES protocol. This ensures that the CE has the complete and up-to-date information of the underlying topology of the Inter-FE network.

The Hello message contains information necessary for discovering and maintaining neighbor relationships. It contains the PE ID, type of protocol element (i.e. CE or FE), interval between any two messages, interval for deeming a neighbor inactive, capability information etc. This is, in some ways, similar to the capabilities of the OSPF Hello protocol.

On receiving the Hello messages from a neighbor, the FE responds back with its own Hello message in a packet format similar to the one received from the neighbor. Essentially, both sides are independently sending Hello messages to each other and listing their neighbor table. Also, each neighbor will see itself listed on its neighbors Hello message. This ensures bi-directionality of the link between any two neighbors.

The operation is concisely described by the following steps:

- . CE activates the topology LFB/component on the FE to initialize on specific ports
- . FE topology LFB/component sends neighbor probes/hellos
- . CE queries FE for its neighbors
- . FE continues to send these probes afterwards (maintenance) and updates asynchronously any new updates

3.2.1. Topology Discovery and Maintenance

Since the CE needs to maintain consistent and up-to-date view of the inter-FE topology, it needs to obtain real-time information of the status of the internal links connecting the FEs. Since the topology

discovery and maintenance occurs during the post-association phase, we make use of the event-notification and query/response messages [ForCESP] of the ForCES protocol to provide this information to the CE. It is important to note that each FE only maintains partial topology information obtained through neighbor relationship maintenance through Hello messages. The partial topology view seen by each FE is only the neighbor connectivity information. The CE has to derive the complete topology view of the interconnected FEs based on the partial topology information reported by each FE (or queried

Ansari et al. Expires: Jan 2005 [Page 7]

Internet Draft ForCES Discovery July 2004

by the CE). This ensures that that only the CE maintains all the intelligence and the protocol operation on the FEs is very simple and has minimal overhead. However, as mentioned above, if optimizations can be performed by having the complete topology information available at the FEs, the CE can push such information to any FE interested in it (interest on the FE may be shown in the form of policy configuration). This is an optional feature available on each FE, which can be turned on or off through configuration or during capability exchange negotiation at setup time. Each FE vendor may decide to make use of this feature in different ways, so the capability to obtain such topology information should exist.

The periodic Hello messages maintain PE neighbor relationships. Any change in the link or neighbor status causes the FE to generate an asynchronous/event-driven message to the CE indicating this change. The mechanism defined in [ForCESP] is used for delivering event-driven messages from the FE to the CE. This involves the CE subscribing to such event-driven messages from the FE.

The CE aggregates the partial topology information received from each FE and generates the inter-FE topology. With this complete knowledge of the inter-FE topology, it can now make appropriate updates to the LFB tables on each FE to move packets inside the NE ; from ingress FE to egress FE, assuming that the destination of the packet is not the current NE itself. Any changes in the internal link states (and hence the topology) requires that the CE reconfigure the LFB tables on the FEs based on the most up-to-date information to ensure that the packets do not end up in a black hole or enter a loop.

3.2.2. Full topology computation at the CE from partial topologies

The CE receives neighbor relationships information from each FE that it uses to construct the full topology of the internal network. Each FE's neighbor relationship table contains information regarding the local element ID, local port connecting the neighbor, the neighbor's

ID, the neighbor's port and any optional additional information. Note that the fact that the FE already knows the neighbor's port information implies that it received the probe/hello messages from the neighbor on that port in response to the hello sent and was, therefore, able to establish bi-directionality of the link. If all the links in the internal network are point-to-point links, the CE simply has to aggregate all the neighbor relationship tables obtained from all the FEs to generate the full topology. If we assume the topology to be a graph, each edge of the graph will be present twice ; essentially providing the same information from the two endpoints of the graph. After deleting all the duplicate entries (and thus reducing the table size by half), the CE now has accurate

Ansari et al. Expires: Jan 2005 [Page 8]

Internet Draft ForCES Discovery July 2004
 view of the full topology. Please refer [section 3.4](#) [Fig. 3(b)] for more details.

[Sub-section on generating full topology from partial topology information for broadcast/multi-access, point-to-multipoint etc. type of links]

3.3. Protocol and Message Headers

The protocol message consists of a fixed length header (16 bytes) followed by one or more optional TLVs. The format of the message is as follows.

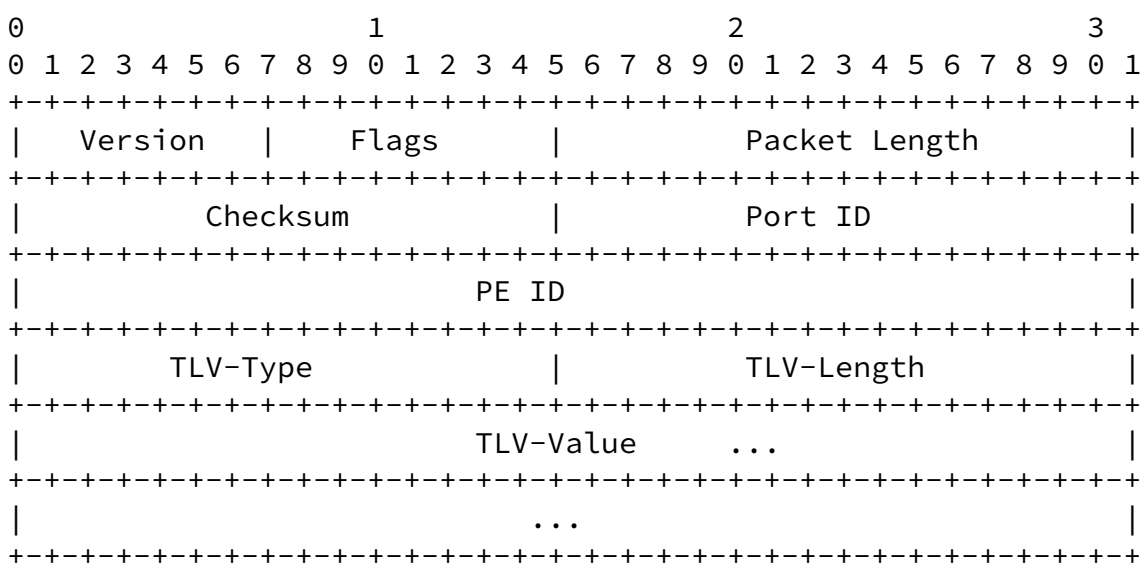


Figure (2)

Version: Version number of this protocol. Currently acceptable value is 0x01

Flags: These indicate whether the message is sent by a FE, (0x01) or CE (0x02). More options may be defined in the future.

Packet Length: The length of the protocol message in bytes, including the header and the following TLVs.

Checksum: 16-bit checksum for the protocol message. The checksum calculation does not include the IP header.

Port ID: This indicates the port on which this packet was sent out by the sender ; useful for topology construction.

PE ID: This is the 32-bit identifier of the sender. It could either be CE ID or FE ID, depending on the sender.

The protocol header is followed by one or many TLVs. The following TLVs types are defined:

Hello TLV: Indicates the Hello message as exchanged by the neighbors. The TLV defines the common hello parameters such as the Hello Interval, Hold time, Uni-directional targeted Hellos, Sequence space number, if needed etc.

Capabilities TLV: Provides the capabilities information - TBD


Vendor specific TLV: TBD


[3.3.1.](#) TLV definitions

TBD

[3.4.](#) Inter-FE Topology Discovery Examples

The following examples illustrate the topology discovery mechanism. For sake of simplicity, we assume that there is only one CE per NE. The FEIDs of the FEs in the topologies below are FE1, FE2, FE3, and FE4. Each FE has port IDs labeled alphabetically. This is also the case with the CE.





FE3 Control Element reachability Table

| <Dest Addr> | <local intf> |
|-------------|--------------|
| CE | A |

FE3 NEIGHBOR ASSOCIATION TABLE

| <local intf> | <neighbor_FEID> | <neighbor_portID> |
|--------------|-----------------|-------------------|
| B | FE2 | E |
| C | FE1 | B |


Figure 3. (a) Full mesh among FE1, FE2, and FE3

During the element-binding phase, each FE sends out hello messages with its FEID and Port ID (as outlined earlier) to all of its neighbors. Since each neighboring FE also listens to such messages, it receives the hello message and adds it to the neighbor association table, which may look like that shown in Fig.4(a). In the topology discovery phase, which is post ForCES association stage, the CE queries each FE about its neighbor table. The FE responds back with the partial topology information available through its neighbor relationships. Both the query and the response are carried by the ForCES protocol. The CE collects the partial topology information from all the FEs in the NE and aggregates this information to fully construct the inter-FE topology. Any changes to the FE neighbor table, e.g. when a link state changes, generates a trigger/update message to the CE. The new information is used to recalculate the new topology and subsequently the CE takes appropriate actions based on the new topology ; such as updating the

packet forwarding tables on the FEs.

The following examples show the neighbor association tables.

3.4.1. Forwarding Elements connected in a daisy chain



Ansari et al.


Expires: Jan 2005

[Page 11]

Internet Draft

ForCES Discovery

July 2004



FE1 NBR ASSOCIATION TABLE

```

-----
<locl intf> <nbr_FEID> <nbr_port>
      B           FE2           E
  
```

FE2 NBR ASSOCIATION TABLE

```

-----
<locl intf> <nbr_FEID> <nbr_port>
      E           FE1           B
      B           FE3           E
  
```

FE3 NBR ASSOCIATION TABLE

```

-----
<locl intf> <nbr_FEID> <nbr_port>
      B           FE4           D
      E           FE2           B
  
```

FE4 NBR ASSOCIATION TABLE

```

-----
<locl intf> <nbr_FEID> <nbr_port>
      D           FE3           B
  
```

CE Topology (Aggregate)View

```

-----
<Node> <Port> <Node> <Port>
  FE1    B    FE2    E\
  FE2    E    FE1    B/
  FE2    B    FE3    E\
  FE3    E    FE2    B/
  
```

CE Topology View

```

-----
<Node> <Port> <Node> <Port>
  FE1    B    FE2    E
  FE2    B    FE3    E
  FE3    B    FE4    D
  -----
  
```

=>


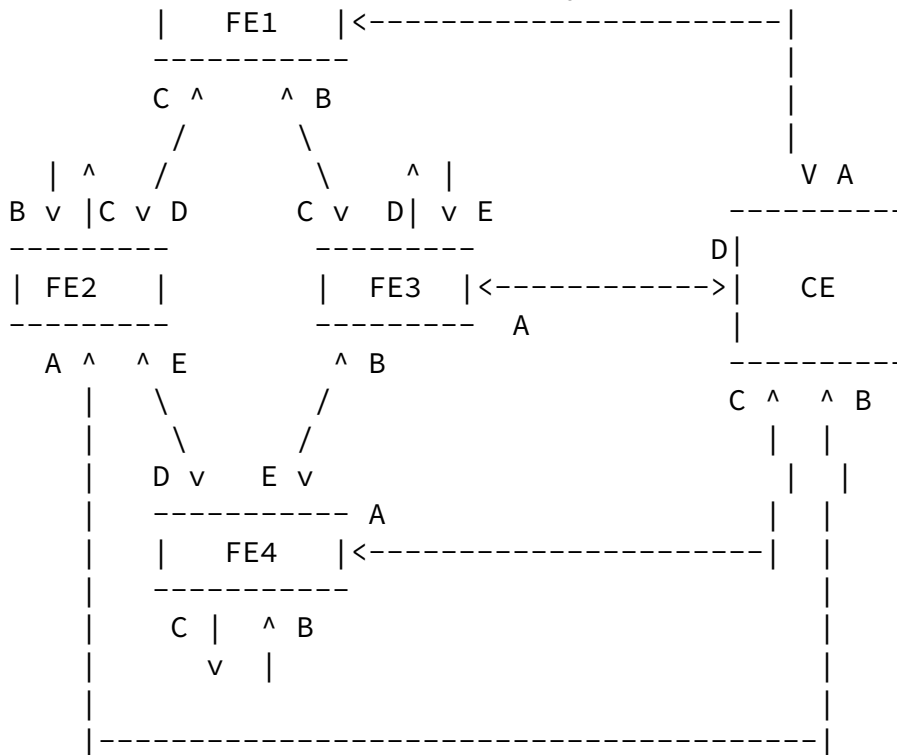


Fig.3(b) Multiple FEs in a daisy chain

3.4.2. Forwarding Elements connected in a ring



FE1 NBR ASSOCIATION TABLE

| <locl_intf> | <nbr_FEID> | <nbr_port> |
|-------------|------------|------------|
| B | FE3 | C |
| C | FE2 | D |

FE2 NBR ASSOCIATION TABLE

| <locl_intf> | <nbr_FEID> | <nbr_port> |
|-------------|------------|------------|
| E | FE4 | D |
| D | FE1 | C |

| FE3 NBR ASSOCIATION TABLE | | | FE4 NBR ASSOCIATION TABLE | | |
|---------------------------|------------|------------|---------------------------|------------|------------|
| <locl_intf> | <nbr_FEID> | <nbr_port> | <locl_intf> | <nbr_FEID> | <nbr_port> |
| B | FE4 | E | D | FE2 | E |
| C | FE1 | B | E | FE3 | B |

Fig. 3(c) Multiple FEs connected in a ring

4. Security Considerations

Like all protocols, this protocol will have security issues as well. These issues will be researched in detail in future draft versions.

5. References

5.1. Normative

Ansari et al.

Expires: Jan 2005

[Page 13]

Internet Draft

ForCES Discovery

July 2004

- [RFC3746] Yang, L., Dantu, R., Anderson, T. and R. Gopal, "Forwarding and Control Element Separation (ForCES) Framework", [RFC 3746](#), April 2004.
- [[RFC3654](#)] Khosravi, H. and T. Anderson, "Requirements for Separation of IP Control and Forwarding", [RFC 3654](#), November 2003.
- [ForCESP] F P Team, "ForCES protocol specification", [draft-ietf-forces-protocol-00.txt](#), Sept 2004.

5.2. Informative

- [OSPF] J. Moy, "OSPF Version 2", 1998, [RFC 2328](#).
- [BGP] Y. Rekhter, T. Li, "A Border Gateway Protocol 4 (BGP-4)", 1995, [RFC 1771](#).
- [IS-IS] R. Collela et al., "Guidelines for OSI NSAP Allocation in the Internet", 1994, [RFC 1629](#).

6. Authors' Addresses

Furquan Ansari
 Bell Labs Research, Lucent Tech.
 101 Crawfords Corner Road
 Holmdel, NJ 07733
 USA
 Phone: +1 732-949-5249

Email: furquan@lucent.com

Hormuzd Khosravi
Intel
2111 N.E. 25th Avenue JF3-206
Hillsboro, OR 97124-5961
USA
Phone: +1 503 264 0334
Email: hormuzd.m.khosravi@intel.com

Jamal Hadi Salim
ZNYX Networks
Ottawa, Ontario, Canada
Email: hadi@znyx.com

7. Full Copyright Notice

Ansari et al.

Expires: Jan 2005

[Page 14]

Internet Draft

ForCES Discovery

July 2004

"Copyright (C) The Internet Society (year). This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights."

"This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE."

8. Acknowledgements

We would like to thank Thyaga Nandagopal of Lucent Technologies for his thoughts and contributions to the initial draft.

Funding for the RFC Editor function is currently provided by the Internet Society.

Ansari et al.

Expires: Jan 2005

[Page 15]