

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: July 12, 2012

W. Wang
Zhejiang Gongshang University
K. Ogawa
NTT Corporation
E. Haleplidis
University of Patras
M. Gao
Hangzhou BAUD Networks
J. Hadi Salim
Mojatatu Networks
January 9, 2012

**Interoperability Report for Forwarding and Control Element Separation
(ForCES)
draft-ietf-forces-interop-03**

Abstract

This document captures test results from the second Forwarding and control Element Separation (ForCES) interoperability test which took place on February 24-25, 2011 in the Internet Technology Lab (ITL) of Zhejiang Gongshang University, China.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on July 12, 2012.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents

(<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
1.1.	ForCES Protocol	3
1.2.	ForCES FE Model	3
1.3.	Transport Mapping Layer	3
2.	Terminology and Conventions	5
2.1.	Requirements Language	5
2.2.	Definitions	5
3.	Overview	7
3.1.	Date, Location, and Participants	7
3.2.	Testbed Configuration	7
3.2.1.	Participants Access	7
3.2.2.	Testbed Configuration	8
4.	Scenarios	11
4.1.	Scenario 1 - LFB Operation	11
4.2.	Scenario 2 - TML with IPSec	11
4.3.	Scenario 3 - CE High Availability	12
4.4.	Scenario 4 - Packet forwarding	14
5.	Test Results	17
5.1.	LFB Operation Test	17
5.2.	TML with IPSec Test	23
5.3.	CE High Availability Test	23
5.4.	Packet Forwarding Test	24
6.	Discussions	27
6.1.	On Data Encapsulation Format	27
7.	Contributors	30
8.	Acknowledgements	31
9.	IANA Considerations	32
10.	Security Considerations	33
11.	References	34
11.1.	Normative References	34
11.2.	Informative References	34
	Authors' Addresses	36

1. Introduction

This document captures the results of the second interoperability test of the Forwarding and control Element Separation (ForCES) Framework which took place February 24-25, 2011 in the Internet Technology Lab (ITL) of Zhejiang Gongshang University, China. The test involved several documents namely: ForCES protocol [[RFC5810](#)] , ForCES FE model [[RFC5812](#)] , ForCES TML [[RFC5811](#)] , ForCES LFB Library [[I-D.ietf-forces-lfb-lib](#)] and ForCES CE HA specification [[I-D.ietf-forces-ceha](#)]. Three independent ForCES implementations participated in the test.

Scenarios of ForCES LFB Operation, TML with IPSec, CE High Availability, and Packet Forwarding are constructed. Series of testing items for every scenario are carried out and interoperability results are achieved. Popular packet analyzers Ethereal/Wireshark[Ethereal] and Tcpdump[Tcpdump] are used to verify the wire results.

The first interoperability test on ForCES was held in July 2008 at the University of Patras, Greece. The test focused on validating the basic semantics of the ForCES protocol and ForCES FE model. The test results were captured by [[RFC6053](#)] .

1.1. ForCES Protocol

The ForCES protocol works in a master-slave mode in which FEs are slaves and CEs are masters. The protocol includes commands for transport of Logical Function Block (LFB) configuration information, association setup, status, and event notifications, etc. The reader is encouraged to read the ForCES protocol specification [[RFC5810](#)] for further information.

1.2. ForCES FE Model

The ForCES FE model [[RFC5812](#)] presents a formal way to define FE Logical Function Blocks (LFBs) using XML. LFB configuration components, capabilities, and associated events are defined when the LFB is formally created. The LFBs within the FE are accordingly controlled in a standardized way by the ForCES protocol.

1.3. Transport Mapping Layer

The ForCES Transport Mapping Layer (TML) transports the ForCES Protocol Layer (PL) messages. The TML is where the issues of how to achieve transport level reliability, congestion control, multicast, ordering, etc are handled. It is expected that more than one TML will be standardized. The various possible TMLs could vary their

implementations based on the capabilities of underlying media and transport. However, since each TML is standardized, interoperability is guaranteed as long as both endpoints support the same TML. All ForCES Protocol Layer implementations MUST be portable across all TMLs. Although more than one TML may be standardized for the ForCES Protocol, for the purposes of the interoperability test, the mandated MUST IMPLEMENT SCTP TML [[RFC5811](#)] was be used.

2. Terminology and Conventions

2.1. Requirements Language

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 [[RFC2119](#)].

2.2. Definitions

This document follows the terminology defined by ForCES related documents, including [RFC3654](#), [RFC3746](#), [RFC5810](#), [RFC5811](#), [RFC5812](#), [RFC5813](#), etc. Some definitions are repeated below for clarity.

Control Element (CE) - A logical entity that implements the ForCES protocol and uses it to instruct one or more FEs on how to process packets. CEs handle functionality such as the execution of control and signaling protocols.

Forwarding Element (FE) - A logical entity that implements the ForCES protocol. FEs use the underlying hardware to provide per-packet processing and handling as directed/controlled by one or more CEs via the ForCES protocol.

LFB (Logical Functional Block) - The basic building block that is operated on by the ForCES protocol. The LFB is a well defined, logically separable functional block that resides in an FE and is controlled by the CE via the ForCES protocol. The LFB may reside at the FE's datapath and process packets or may be purely an FE control or configuration entity that is operated on by the CE. Note that the LFB is a functionally accurate abstraction of the FE's processing capabilities, but not a hardware-accurate representation of the FE implementation.

LFB Class and LFB Instance - LFBs are categorized by LFB Classes. An LFB Instance represents an LFB Class (or Type) existence. There may be multiple instances of the same LFB Class (or Type) in an FE. An LFB Class is represented by an LFB Class ID, and an LFB Instance is represented by an LFB Instance ID. As a result, an LFB Class ID associated with an LFB Instance ID uniquely specifies an LFB existence.

LFB Metadata - Metadata is used to communicate per-packet state from one LFB to another, but is not sent across the network. The FE model defines how such metadata is identified, produced, and consumed by the LFBs. It defines the functionality but not how metadata is encoded within an implementation.

LFB Components - Operational parameters of the LFBs that must be visible to the CEs are conceptualized in the FE model as the LFB components. The LFB components include, for example, flags, single-parameter arguments, complex arguments, and tables that the CE can read and/or write via the ForCES protocol.

ForCES Protocol - While there may be multiple protocols used within the overall ForCES architecture, the term "ForCES protocol" and "protocol" refer to the "Fp" reference points in the ForCES framework in [[RFC3746](#)]. This protocol does not apply to CE-to-CE communication, FE-to-FE communication, or to communication between FE and CE managers. Basically, the ForCES protocol works in a master-slave mode in which FEs are slaves and CEs are masters.

ForCES Protocol Transport Mapping Layer (ForCES TML) - A layer in ForCES protocol architecture that uses the capabilities of existing transport protocols to specifically address protocol message transportation issues, such as how the protocol messages are mapped to different transport media (like TCP, IP, ATM, Ethernet, etc.), and how to achieve and implement reliability, multicast, ordering, etc. The ForCES TML specifications are detailed in separate ForCES documents, one for each TML.

3. Overview

3.1. Date, Location, and Participants

The second ForCES interoperability test meeting was held by IETF ForCES Working Group on February 24-25, 2011, and was chaired by Jamal Hadi Salim. Three independent ForCES implementations participated in the test:

- * Zhejiang Gongshang University/Hangzhou BAUD Corporation of Information and Networks Technology (Hangzhou BAUD Networks), China. This implementation is referred to as "China" or in some cases "C" in the document for the sake of brevity.
- * NTT Corporation, Japan. This implementation is referred to as "Japan" or in some cases "J" in the document for the sake of brevity.
- * The University of Patras, Greece. This implementation is referred to as "Greece" or in some cases "G" in the document for the sake of brevity.

Two other organizations, Mojatatu Networks and Hangzhou BAUD Networks Corporation, which independently extended two different well known public domain protocol analyzers, Ethereal/Wireshark [[Ethereal](#)] and Tcpdump [[Tcpdump](#)], also participated in the interop test. During the interoperability test, the two protocol analyzers were used to verify the validity of ForCES protocol messages and in some cases semantics.

Some issues related to interoperability among implementations were discovered. Most of the issues were solved on site during the test. The most contentious issue found was on the format of encapsulation for protocol TLV (Refer to [Section 6.1](#)).

Some errata related to ForCES document were found by the interoperability test. The errata has been reported to related IETF RFCs.

At times, interoperability testing was exercised between two instead of all three representative implementations due to a third one lacking a specific feature; however, in ensuing discussions, all implementers mentioned they will be implementing any missing features in the future.

3.2. Testbed Configuration

3.2.1. Participants Access

Japan and China physically attended on site at the Internet Technology Lab (ITL) of Zhejiang Gongshang University in China. The University of Patras implementation joined remotely from Greece. The

chair, Jamal Hadi Salim, joined remotely from Canada by using the Teamviewer as the monitoring tool[Teamviewer]. The approach is as shown in Figure 1. In the figure, FE/CE refers to FE or CE that the implementer may act alternatively.




Figure 1: Access for Participants

As specified in [RFC 5811](#), all CEs and FEs SHALL implement IPSec security in the TML.

On the internet boundary, gateways used MUST allow for IPSec, SCTP protocol and SCTP ports as defined in the ForCES SCTP-TML [[RFC5811](#)] .

3.2.2. Testbed Configuration

CEs and FEs from China and Japan implementations were physically located within the ITL Lab of Zhejiang Gongshang University and connected together using Ethernet switches. The configuration can be seen in Figure 2. In the figure, the SmartBits is a third-party supplied routing protocol testing machine, which acts as a router running OSPF and RIP and exchanges routing protocol messages with ForCES routers in the network. The Internet is connected via an ADSL channel.

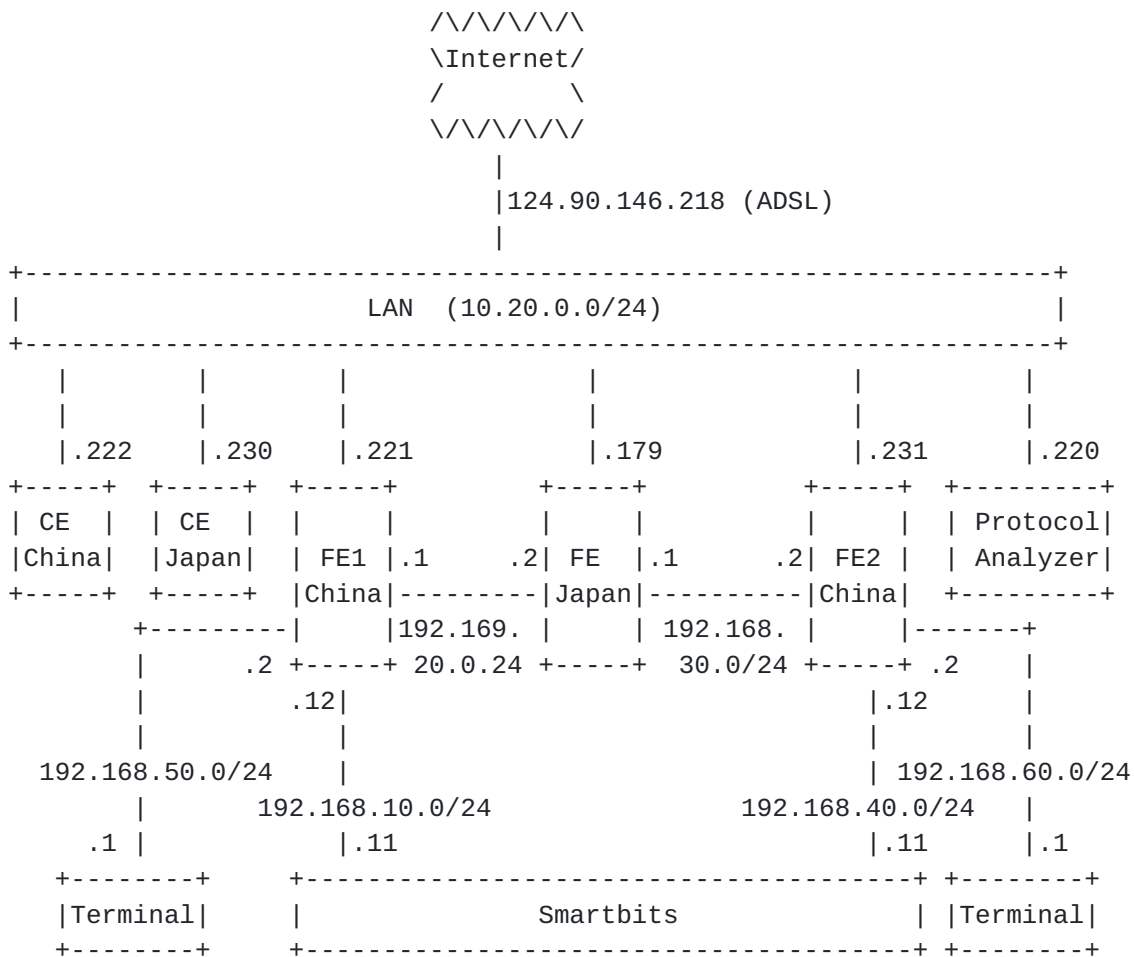


Figure 2: Testbed Configuration Located in ITL Lab,China

Hardwares and Softwares (CE and FE) of Greece that were located within the University of Patras, Greece, were connected together using LAN as shown in Figure 3. The Internet is connected via a VPN channel.

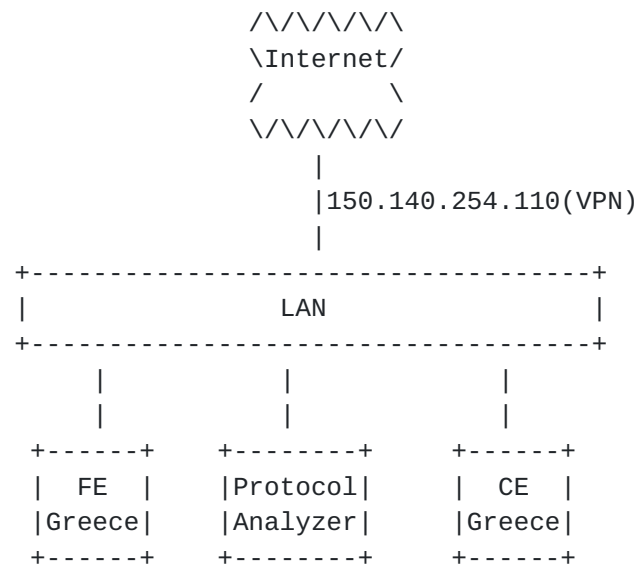


Figure 3: Testbed Configuration Located in the University of Patras, Greece

All above testbed configurations can then satisfy requirements of all the interoperability test scenarios that are mentioned in this document.

4. Scenarios

4.1. Scenario 1 - LFB Operation

This scenario is to test the interoperability on LFB operations among the participants. The connection diagram for the participants is as shown in Figure 4.

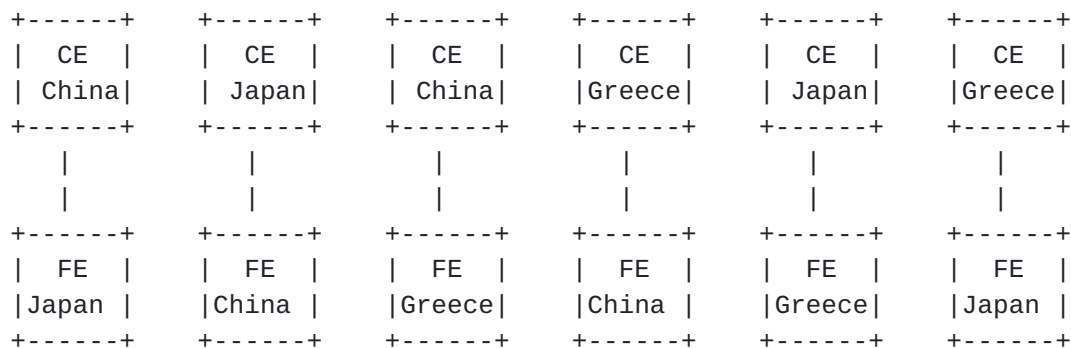


Figure 4: Scenario for LFB Operation

In order to make interoperability more credible, the three implementers are required to carry out the test in a way acting as CE or FE alternatively. As a result, every LFB operation is combined with 6 scenarios, as shown by Figure 4.

The test scenario is designed with the following purposes:

Firstly, the scenario is designed to verify all kinds of protocol messages with their complex data formats, which are defined in [RFC 5810](#). Specially, we try to verify the data format of a PATH-DATA with nested PATH-DATAs, and the operation(SET, GET, DEL) of an array or an array with a nested array.

Secondly, the scenario is designed to verify the definition of ForCES LFB Library[FORCES-LFBLIB], which defines a base set of ForCES LFB classes for typical router functions. Successful test under this scenario also means the validity of the LFB definitions.

4.2. Scenario 2 - TML with IPSec

This scenario is designed to implement a TML with IPSec, which is the requirement by [RFC 5811](#). TML with IPSec was not implemented in the first ForCES interoperability test as reported by [RFC 6053](#). For this reason, in the second interoperability test, we specifically designed the test scenario to verify the TML over IPSec channel.

In this scenario, tests on LFB operations for Scenario 1 were

repeated with the difference that TML was secured via IPSec. This setup scenario allows us to verify whether all interactions between CE and FE can be made correctly under an IPSec TML environment.

The connection diagram for this scenario is shown as Figure 5. Because of system deficiency to deploy IPSec over TML in Greece, the text only took place between China and Japan.

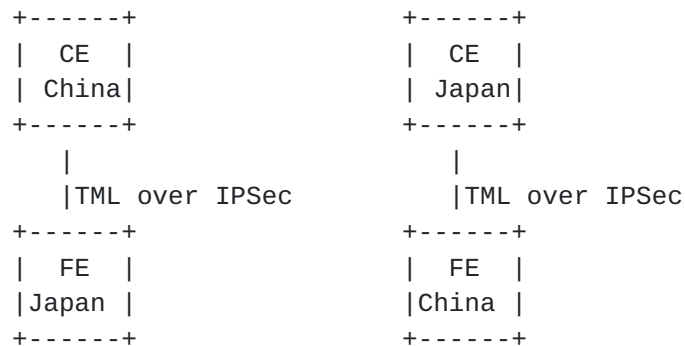


Figure 5: Scenario for LFB Operation with TML over IPSec

In this scenario, ForCES TML was run over IPSec channel. Implementers joined in this interoperability have used the same third-party software 'racoon' to have established the IPSec channel.

China and Japan have made a successful test with the scenario, and the following items have been realized:

- o Internet Key Exchange (IKE) with certificates for endpoint authentication.
- o Transport Mode Encapsulating Security Payload (ESP). HMAC-SHA1-96 [[RFC2404](#)] for message integrity protection.

4.3. Scenario 3 - CE High Availability

CE High Availability (CEHA) was tested based on the ForCES CEHA document [ForCES-CEHA].

The design of the setup and the scenario for the CEHA were simplified so as to focus mostly on the mechanics of the CEHA, which are:

- o Associating with more than one CE.
- o Switching to backup CE on master CE failure.

The connection diagram for the scenario is as shown in Figure 6.

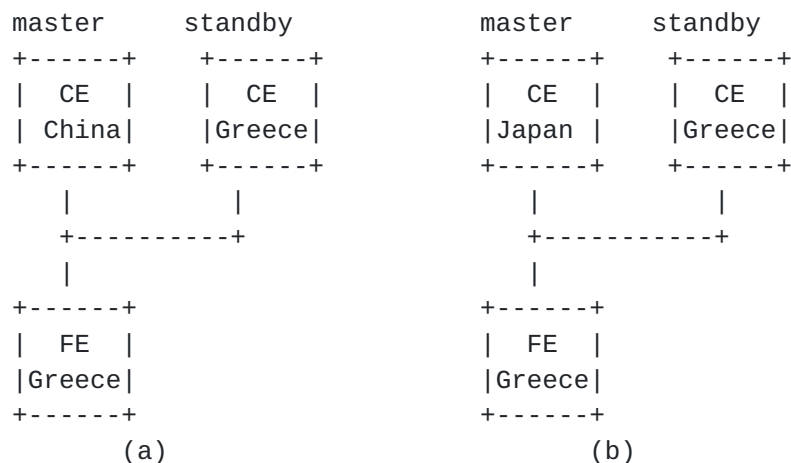


Figure 6: Scenario for CE High Availability

In this scenario one FE is connected and associated to a master CE and a backup CE. In the pre-association phase, the FE would be configured to have China's or Japan's CE as master CE and Greece's CE as standby CE. The CEFailoverPolicy component of the FE Protocol Object LFB that specifies whether the FE is in High Availability mode (value 2 or 3) would either be set in the pre-association phase by the FEM interface or in post-association phase by the master CE.

If the CEFailoverPolicy value is set to 2 or 3, the FE (in the post-association phase) will attempt to connect and associate with the standby CE.

When the master CE is deemed disconnected, either by a TearDown, Loss of Heartbeats or physically disconnected, the FE would assume that the standby CE is now the master CE. The FE will then send an Event Notification, Primary CE Down, to all associated CEs, only the standby CE in this case, with the value of the new master CEID. The standby CE will then respond by sending a configuration message to the CEID component of the FE Protocol Object with its own ID to confirm that the CE considers itself as the master as well.

The steps of the CEHA test scenario are as follows:

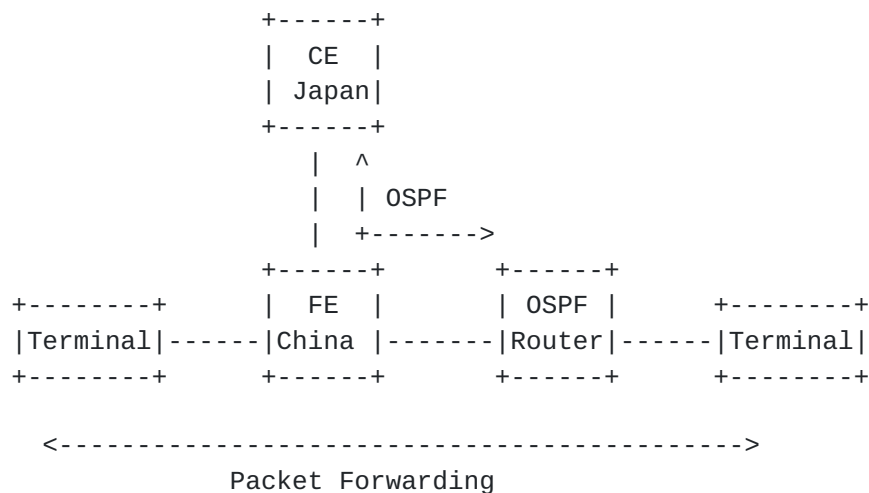
1. In the pre-association phase, setup of FE with master CE and backup CE
2. FE connecting and associating with master CE.
3. When CEFailoverPolicy is set to 2 or 3, the FE will connect and associate with backup CE.

4. Once the master CE is considered disconnected then the FE chooses the first Associated backup CE.
5. It sends an Event Notification specifying that the master CE is down and who is now the master CE.
6. The new master CE sends a SET Configuration message to the FE setting the CEID value to who is now the new master CE completing the switch.

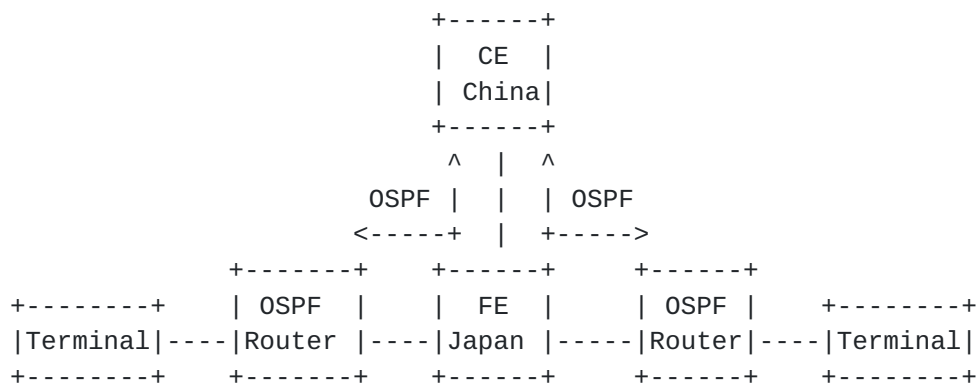
4.4. Scenario 4 - Packet forwarding

This test scenario is to verify LFBs like RedirectIn, RedirectOut, IPv4NextHop, IPv4UcastLPM defined by the ForCES LFB library document[ForCES-LFBLIB], and more importantly, to verify the combination of the LFBs to implement IP packet forwarding.

The connection diagram for this scenario is as Figure 7.



(a)






Figure 7: Scenario for IP Packet forwarding

In case (a), a CE by Japan is connected to an FE by China to form a ForCES router. A Smartbits test machine with its routing protocol software are used to simulate an OSPF router and are connected with the ForCES router to try to exchange OSPF hello packets and LSA packets among them. Terminals are simulated by Smartbits to send and receive packets. As a result, the CE in the ForCES router need to be configured to run and support OSPF routing protocol.

In case (b), a CE by China is connected to an FE by Japan to form a ForCES router. Two routers running OSPF are simulated and connected to the ForCES router to test if the ForCES router can support OSPF protocol and support packet forwarding.

In case (c), two ForCES routers are constructed. One is with CE by Japan and FE by China and the other is opposite. OSPF and packet forwarding are tested in the environment.

Testing process for this scenario is as below:

1. Boot terminals and routers, and set IP addresses of their interfaces.

2. Boot CE and FE.
3. Establish association between CE and FE, and set IP addresses of FEs interfaces.
4. Start OSPF among CE and routers, and set FIB on FE.
5. Send packets between terminals.

5. Test Results

5.1. LFB Operation Test

The test result is as reported by Figure 8. For the convenience sake, as mentioned earlier, abbreviations of 'C' in the table means implementation from China, 'J' Japan implementation, and 'G' Greece implementation.

Test#	CE	FE(s)	Oper	LFB	Component /Capability	Result
1	C	J	GET	FEObject	LFBTopology	Success
	J	C				Success
	G	C				Success
	J	G				Success
	G	J				Success
2	C	J	GET	FEObject	LFBSelector	Success
	J	C				Success
	C	G				Success
	G	C				Success
	J	G				Success
	G	J				Success
3	C	J	GET	EtherPHYCop	PHYPortID	Success
	J	C				Success
	C	G				Success
	G	C				Success
	J	G				Success
	G	J				Success
4	C	J	GET	EtherPHYCop	AdminStatus	Success
	J	C				Success
	C	G				Success
	G	C				Success
	J	G				Success
	G	J				Success
5	C	J	GET	EtherPHYCop	OperStatus	Success
	J	C				Success
	C	G				Success
	G	C				Success
	J	G				Success
	G	J				Success
6	C	J	GET	EtherPHYCop	AdminLinkSpeed	Success

	J	C					Success
	C	G					Success
	G	C					Success
	J	G					Success
	G	J					Success
7	C	J	GET	EtherPHYCop	OperLinkSpeed		Success
	J	C					Success
	C	G					Success
	G	C					Success
	J	G					Success
	G	J					Success
8	C	J	GET	EtherPHYCop	AdminDuplexSpeed		Success
	J	C					Success
	C	G					Success
	G	C					Success
	J	G					Success
	G	J					Success
9	C	J	GET	EtherPHYCop	OperDuplexSpeed		Success
	J	C					Success
	C	G					Success
	G	C					Success
	J	G					Success
	G	J					Success
10	C	J	GET	EtherPHYCop	CarrierStatus		Success
	J	C					Success
	C	G					Success
	G	C					Success
	J	G					Success
	G	J					Success
11	C	J	GET	EtherMACIn	AdminStatus		Success
	J	C					Success
	C	G					Success
	G	C					Success
	J	G					Success
	G	J					Success
12	C	J	GET	EtherMACIn	LocalMacAddresses		Success
	J	C					Success
	C	G					Success
	G	C					Success
	J	G					Success
	G	J					Success

13	C	J	GET	EtherMACIn	L2Bridging PathEnable	Success
	J	C				Success
	C	G				Success
	G	C				Success
	J	G				Success
	G	J				Success
14	C	J	GET	EtherMACIn	PromiscuousMode	Success
	J	C				Success
	C	G				Success
	G	C				Success
	J	G				Success
	G	J				Success
15	C	J	GET	EtherMACIn	TxFlowControl	Success
	J	C				Success
	C	G				Success
	G	C				Success
	J	G				Success
	G	J				Success
16	C	J	GET	EtherMACIn	RxFlowControl	Success
	J	C				Success
	C	G				Success
	G	C				Success
	J	G				Success
	G	J				Success
17	C	J	GET	EtherMACIn	MACInStats	Success
	J	C				Success
	C	G				Success
	G	C				Success
	J	G				Success
	G	J				Success
18	C	J	GET	EtherMACOut	AdminStatus	Success
	J	C				Success
	C	G				Success
	G	C				Success
	J	G				Success
	G	J				Success
19	C	J	GET	EtherMACOut	MTU	Success
	J	C				Success
	C	G				Success
	G	C				Success
	J	G				Success
	G	J				Success

20	C	J	GET	EtherMACOut	TxFlowControl	Success
	J	C				Success
	C	G				Success
	G	C				Success
	J	G				Success
	G	J				Success
21	C	J	GET	EtherMACOut	TxFlowControl	Success
	J	C				Success
	C	G				Success
	G	C				Success
	J	G				Success
	G	J				Success
22	C	J	GET	EtherMACOut	MACOutStats	Success
	J	C				Success
	C	G				Success
	G	C				Success
	J	G				Success
	G	J				Success
23	C	J	GET	ARP	PortV4AddrInfoTable	Success
	J	C				Success
	C	G				Success
	G	C				Success
	J	G				Success
	G	J				Success
24	C	J	SET	ARP	PortV4AddrInfoTable	Success
	J	C				Success
	C	G				Success
	G	C				Success
	J	G				Success
	G	J				Success
25	C	J	DEL	ARP	PortV4AddrInfoTable	Success
	J	C				Success
	C	G				Success
	G	C				Success
	J	G				Success
	G	J				Success
26	C	J	SET	EtherMACIn	LocalMACAddresses	Success
	J	C				Success
	C	G				Success
	G	C				Success
	J	G				Success
	G	J				Success

	G	J					Success
27	C	J	SET	EtherMACIn		MTU	Success
	J	C					Success
	C	G					Success
	G	C					Success
	J	G					Success
	G	J					Success
28	C	J	SET	IPv4NextHop		IPv4NextHopTable	Success
	J	C					Success
	C	G					Success
	G	C					Success
	J	G					Success
	G	J					Success
29	C	J	SET	IPv4UcastLPM		IPv4PrefixTable	Success
	J	C					Success
	C	G					Success
	G	C					Success
	J	G					Success
	G	J					Success
30	C	J	DEL	IPv4NextHop		IPv4NextHopTable	Success
	J	C					Success
	C	G					Success
	G	C					Success
	J	G					Success
	G	J					Success
31	C	J	DEL	IPv4UcastLPM		IPv4PrefixTable	Success
	J	C					Success
	C	G					Success
	G	C					Success
	J	G					Success
	G	J					Success
32	C	J	SET	EtherPHYCop		AdminStatus	Success
	J	C					Success
	C	G					Success
	G	C					Success
	J	G					Success
	G	J					Success
33	C	J	SET	Ether Classifier		VlanInputTable	Success
	J	C					Success
	C	G					Success
	G	C					Success

	J	G					Success
	G	J					Success
34	C	J	DEL	Ether	VlanInputTable	Success	
	J	C		Classifier		Success	
	C	G				Success	
	G	C				Success	
	J	G				Success	
	G	J				Success	
35	C	J	SET	Ether	VlanOutputTable	Success	
	J	C		Encapsulator		Success	
	C	G				Success	
	G	C				Success	
	J	G				Success	
	G	J				Success	
36	C	J	DEL	Ether	VlanOutputTable	Success	
	J	C		Encapsulator		Success	
	C	G				Success	
	G	C				Success	
	J	G				Success	
	G	J				Success	

Figure 8: LFB Operation Test Results

Note on test 1#:

On the wire format of encapsulation on array, only the case of FULLDATA-in-FULLDATA was tested.

In China's implementation, after test 2# CE have to get all LFBs' instance data actively according to the queried component of LFBSelectors.

Note on test 28# and 29#:

Only had new reachable network destination been set, can route entry be added into system.

Note on test 30# and 31#:

Corresponding nexthop entry must be deleted before prefix entry which is decided by FE's routing management.

5.2. TML with IPSec Test

In this scenario, the ForCES TML is run over IPSec. Implementers joined this interoperability test use the same third-party tool software 'racoon' to establish IPSec channel. Some typical LFB operation tests as in Scenario 1 are repeated with the IPSec enabled TML.

A note on this test is, because of the system difficulty to implement IPSec over TML, Greece did not join in the test. Therefore, this scenario only took place between C and J.

The TML with IPSec test results are reported by Figure 9.

Test#	CE	FE(s)	Oper	LFB	Component/ Capability	Result
1	C	J	GET	FEObject	LFBTopology	Success
	J	C				Success
2	C	J	GET	FEObject	LFBSelectors	Success
	J	C				Success
3	C	J	SET	Ether Classifier	VlanInputTable	Success
	J	C				Success
4	C	J	DEL	Ether Classifier	VlanInputTable	Success
	J	C				Success

Figure 9: TML with IPSec Test Results

5.3. CE High Availability Test

In this scenario one FE connects and associates with a master CE and a backup CE. When the master CE is deemed disconnected the FE would attempt to find another associated CE to become the master CE.

The CEHA scenario as is described in Scenario 3 was completed successfully for both setups.

Due to a bug in one of the FEs, a interesting issue was caught: it was observed that the buggy FE took up to a second to failover. It was eventually found that the issue was due to the FE's prioritization of the different CEs. All messages from the backup CE were being ignored unless the master CE is disconnected.

While the bug was fixed and the CEHA scenario was completed successfully, the authors feel it was important to capture the implementation issue in this document. The recommended approach is the following:

- o The FE SHOULD receive and handle messages first from the master CE on all priority channels to maintain proper functionality and then receive and handle messages from the backup CEs.
- o Only when the FE is attempting to associate with the backup CEs, then the FE SHOULD receive and handle messages per priority channel from all CEs. When all backup CEs are associated with or deemed unreachable, then the FE SHOULD return to receiving and handling messages first from the master CE.

5.4. Packet Forwarding Test

As described in the ForCES LFB library [[I-D.ietf-forces-lfb-lib](#)], packet forwarding is implemented by a set of LFB classes that compose a processing path for packets. In this test scenario, as shown in Figure 7, a ForCES router running OSPF protocol was constructed. In addition, a set of LFBs including RedirectIn, RedirectOut, IPv4UcastLPM, and IPv4NextHop LFBs are used. RedirectIn and RedirectOut LFBs redirect OSPF hello and LSA packets from and to CE. A Smartbits test machine is used to simulate an OSPF router and exchange the OSPF hello and LSA packets with CE in ForCES router.

Cases (a) and (b) in Figure 7 both need a RedirectIn LFB to send OSPF packets generated by CE to FE by use of ForCES packet redirect messages. The OSPF packets are further sent to an outside OSPF Router by the FE via forwarding LFBs including IPv4NextHop and IPv4UcastLPM LFBs. A RedirectOut LFB in the FE is used to send OSPF packets received from outside OSPF Router to CE by ForCES packet redirect messages.

By running OSPF, the CE in the ForCES router can generate new routes and load them to routing table in FE. The FE is then able to forward packets according to the routing table.

The test is reported with the results in Figure 10

Test#	CE	FE(s)	Item	LFBs Related	Result
1	J	C	IPv4NextHopTable SET	IPv4NextHop	Success
2	J	C	IPv4PrefixTable SET	IPv4UcastLPM	Success

3	J	C	Redirect ospf packet from CE to SmartBits	RedirectIn	Success
4	J	C	Redirect ospf packet from SmartBits to CE	RedirectOut	Success
5	J	C	Metadata in redirect message	RedirectOut RedirectIn	Success
6	J	C	OSPF neighbor discovery	RedirectOut RedirectIn	Success
7	J	C	OSPF DD exchange	RedirectOut RedirectIn IPv4NextHop	Success
8	J	C	OSPF LSA exchange	RedirectOut RedirectIn IPv4NextHop IPv4UcastLPM	Success
9	J	C	Data Forwarding	RedirectOut RedirectIn IPv4NextHop IPv4UcastLPM	Success
10	C	J	IPv4NextHopTable SET	IPv4NextHop	Success
11	C	J	IPv4PrefixTable SET	IPv4UcastLPM	Success
12	C	J	Redirect ospf packet from CE to other OSPF router	RedirectIn	Success
13	C	J	Redirect ospf packet from other OSPF router to CE	RedirectOut	Success
14	C	J	Metadata in redirect message	RedirectOut RedirectIn	Success
15	C	J	OSPF neighbor discovery	RedirectOut RedirectIn	Success
16	C	J	OSPF DD exchange	RedirectOut RedirectIn IPv4NextHop	Failure
17	C	J	OSPF LSA exchange	RedirectOut RedirectIn	Failure

6. Discussions

6.1. On Data Encapsulation Format

In the first day of the test, it was found that the LFB inter-operations about tables all failed. The reason is found to be the different ForCES protocol data encapsulation method among different implementations. The encapsulation issues are detailed as below:

Assuming that an LFB has two components, one is a struct with ID 1 and the other an array with ID 2, further with two components of u32 both inside, as below:

```
struct1: type struct, ID=1
    components are:
    a, type u32, ID=1
    b, type u32, ID=2
```

```
table1: type array, ID=2
    components for each row are (a struct of):
    x, type u32, ID=1
    y, type u32, ID=2
```

1. On response of PATH-DATA format

When a CE sends a config/query ForCES protocol message to an FE from a different implementer, the CE probably receives response from the FE with different PATH-DATA encapsulation format. For example, if a CE sends a query message with a path of 1 to a third party FE to manipulate struct 1 as defined above, the FE is probable to generate response with two different PATH-DATA encapsulation format: one is the value with FULL/SPARSE-DATA and the other is the value with many parallel PATH-DATA TLV and nested PATH-DATA TLV, as below:

format 1:

```
    OPER = GET-RESPONSE-TLV
    PATH-DATA-TLV:
        IDs=1
        FULLDATA-TLV containing valueof(a),valueof(b)
```

format 2:

```
    OPER = GET-RESPONSE-TLV
    PATH-DATA-TLV:
        IDs=1
        PATH-DATA-TLV:
            IDs=1
            FULLDATA-TLV containing valueof(a)
        PATH-DATA-TLV:
            IDs=2
```


FULLDATA-TLV containing valueof(b)

The interoperability test witnessed that a ForCES element (CE or FE) sender is free to choose whatever data structure that IETF ForCES documents define and best suits the element, while a ForCES element (CE or FE) MUST be able to accept and process information (requests and responses) that use any legitimate structure defined by IETF ForCES documents. While in the case a ForCES element is free to choose any legitimate data structure as a response, it is preferred the ForCES element responds in the same format that the request was made, as it is most probably the data structure is the request sender looks forward to receive.

2. On operation to array

An array operation may also have several different data encapsulation formats. For instance, if a CE sends a config message to table 1 with a path of (2.1), which refers to component with ID=2, which is an array, and the second ID is the row, so row 1, it may be encapsulated with three formats as below:

format 1:

```
OPER = SET-TLV
  PATH-DATA-TLV:
    IDs=2.1
    FULLDATA-TLV containing valueof(x),valueof(y)
```

format 2:

```
OPER = SET-TLV
  PATH-DATA-TLV:
    IDs=2.1
    PATH-DATA-TLV:
      IDs=1
      FULLDATA-TLV containing valueof(x)
    PATH-DATA-TLV
      IDs=2
      FULLDATA-TLV containing valueof(y)
```

Moreover, if CE is targeting the whole array, for example if the array is empty and CE wants to add the first row to the table, it could also adopt another format:

format 3:

```
OPER = SET-TLV
  PATH-DATA-TLV:
    IDs=2
    FULLDATA-TLV containing rowindex=1,valueof(x),valueof(y)
```


The interoperability test experience shows that format 1 and format 3, which take full advantage of multiple data elements description in one TLV of FULLDATA-TLV, get more efficiency, although format 2 can also get the same operating goal.

7. Contributors

Contributors who have made major contributions to the interoperability test are as below:

Hirofumi Yamazaki
NTT Corporation
Tokyo
Japan
Email: yamazaki.horofumi@lab.ntt.co.jp

Rong Jin
Zhejiang Gongshang University
Hangzhou
P.R.China
Email: jinrong@zjgsu.edu.cn

Yuta Watanabe
NTT Corporation
Tokyo
Japan
Email: yuta.watanabe@ntt-at.co.jp

Xiaochun Wu
Zhejiang Gongshang University
Hangzhou
P.R.China
Email: spring-403@zjgsu.edu.cn

8. Acknowledgements

The authors would also like thank the following test participants:

Chuanhuang Li, Hangzhou BAUD Networks
Ligang Dong, Zhejiang Gongshang University
Jingjing Zhou, Zhejiang Gongshang University
Liaoyuan Ke, Hangzhou BAUD Networks
Kelei Jin, Hangzhou BAUD Networks

9. IANA Considerations

This memo includes no request to IANA.

10. Security Considerations

Developers of ForCES FEs and CEs must take the security considerations of the ForCES Framework [[RFC3746](#)] and the ForCES Protocol [[RFC5810](#)] into account. Also, as specified in the security considerations section of the SCTP-Based TML for the ForCES Protocol [[RFC5811](#)] the transport-level security, has to be ensured by IPsec.

11. References

11.1. Normative References

- [RFC5810] Doria, A., Hadi Salim, J., Haas, R., Khosravi, H., Wang, W., Dong, L., Gopal, R., and J. Halpern, "Forwarding and Control Element Separation (ForCES) Protocol Specification", [RFC 5810](#), March 2010.
- [RFC5811] Hadi Salim, J. and K. Ogawa, "SCTP-Based Transport Mapping Layer (TML) for the Forwarding and Control Element Separation (ForCES) Protocol", [RFC 5811](#), March 2010.
- [RFC5812] Halpern, J. and J. Hadi Salim, "Forwarding and Control Element Separation (ForCES) Forwarding Element Model", [RFC 5812](#), March 2010.
- [RFC5813] Haas, R., "Forwarding and Control Element Separation (ForCES) MIB", [RFC 5813](#), March 2010.

11.2. Informative References

- [Ethereal]
"Ethereal, also named Wireshark, is a protocol analyzer. The specific Ethereal that was used is an updated Ethereal, by Fenggen Jia, that can analyze and decode the ForCES protocol messages", <http://www.ietf.org/mail-archive/web/forces/current/msg03687.html> .
- [I-D.ietf-forces-ceha]
Ogawa, K., Wang, W., Haleplidis, E., and J. Salim, "ForCES Intra-NE High Availability", [draft-ietf-forces-ceha-01](#) (work in progress), February 2011.
- [I-D.ietf-forces-lfb-lib]
Wang, W., Haleplidis, E., Ogawa, K., Li, C., and J. Halpern, "ForCES Logical Function Block (LFB) Library", [draft-ietf-forces-lfb-lib-05](#) (work in progress), July 2011.
- [RFC2629] Rose, M., "Writing I-Ds and RFCs using XML", [RFC 2629](#), June 1999.
- [RFC3654] Khosravi, H. and T. Anderson, "Requirements for Separation of IP Control and Forwarding", [RFC 3654](#), November 2003.
- [RFC3746] Yang, L., Dantu, R., Anderson, T., and R. Gopal, "Forwarding and Control Element Separation (ForCES)

Framework", [RFC 3746](#), April 2004.

[RFC6053] Haleplidis, E., Ogawa, K., Wang, W., and J. Hadi Salim, "Implementation Report for Forwarding and Control Element Separation (ForCES)", [RFC 6053](#), November 2010.

[Tcpdump] "Tcpdump is a Linux protocol analyzer. The specific tcpdump that was used is a modified tcpdump, by Jamal Hadi Salim, that can analyze and decode the ForCES protocol messages", <http://www.ietf.org/mail-archive/web/forces/current/msg03811.html> .

[Teamviewer] "TeamViewer connects to any PC or server around the world within a few seconds.", <http://www.teamviewer.com/> .

Authors' Addresses

Weiming Wang
Zhejiang Gongshang University
18 Xuezheng Str., Xiasha University Town
Hangzhou, 310018
P.R.China

Phone: +86-571-28877721
Email: wmwang@zjgsu.edu.cn

Kentaro Ogawa
NTT Corporation
Tokyo,
Japan

Email: ogawa.kentaro@lab.ntt.co.jp

Evangelos Haleplidis
University of Patras
Patras,
Greece

Email: ehalep@ece.upatras.gr

Ming Gao
Hangzhou BAUD Networks
408 Wen-San Road
Hangzhou, 310012
P.R.China

Phone: +86-571-28877751
Email: gmyyqno1@pop.zjgsu.edu.cn

Jamal Hadi Salim
Mojatatu Networks
Ottawa
Canada

Email: hadi@mojatatu.com

