

Internet Engineering Task Force
Internet-Draft
Updates: [6053](#) (if approved)
Intended status: Informational
Expires: December 04, 2013

W. Wang
Zhejiang Gongshang University
K. Ogawa
NTT Corporation
E.H. Haleplidis
University of Patras
M. Gao
Hangzhou BAUD Networks
J. Hadi Salim
Mojatatu Networks
June 02, 2013

**Interoperability Report for Forwarding and Control Element Separation
(ForCES)
draft-ietf-forces-interop-09**

Abstract

This document captures results of 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. [RFC 6053](#) reported the results of the first ForCES interoperability test, and this document updates [RFC 6053](#) by providing further interoperability results.

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 December 04, 2013.

Copyright Notice

Copyright (c) 2013 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	4
1.3.	Transport Mapping Layer	4
1.4.	Definitions	4
2.	Overview	4
2.1.	Date, Location, and Participants	4
2.2.	Testbed Configuration	5
2.2.1.	Participants Access	5
2.2.2.	Testbed Configuration	6
3.	Scenarios	7
3.1.	Scenario 1 - LFB Operation	7
3.2.	Scenario 2 - TML with IPsec	8
3.3.	Scenario 3 - CE High Availability	9
3.4.	Scenario 4 - Packet forwarding	11
4.	Test Results	13
4.1.	LFB Operation Test	13
4.2.	TML with IPsec Test	19
4.3.	CE High Availability Test	19
4.4.	Packet Forwarding Test	20
5.	Discussions	22
5.1.	On Data Encapsulation Format	22
6.	Contributors	25
7.	Acknowledgements	25
8.	IANA Considerations	26
9.	Security Considerations	26
10.	References	26
10.1.	Normative References	26
10.2.	Informative References	26
	Authors' Addresses	28

1. Introduction

This document captures results of the second interoperability test of the Forwarding and Control Element Separation (ForCES) which took place February 24-25, 2011 in the Internet Technology Lab (ITL) of Zhejiang Gongshang University, China. The test involved protocol elements described in several documents namely:

- ForCES Protocol [[RFC5810](#)]
- ForCES Forwarding Element (FE) Model [[RFC5812](#)]
- ForCES Transport Mapping Layer (TML) [[RFC5811](#)]

The test also involved protocol elements described in the then-current versions of two Internet-Drafts. Although these documents have subsequently been revised and advanced, it is important to understand which versions of the work were used during this test. The then-current Internet-Drafts are:

- ForCES Logical Function Block (LFB) Library [[I-D.ietf-forces-lfb-lib-03](#)].
- ForCES Intra-NE High Availability [[I-D.ietf-forces-ceha-00](#)].

The 'ForCES Logical Function Block (LFB) Library' document has advanced and been published by IETF as [RFC 6956](#).

Three independent ForCES implementations participated in the test.

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

This document is an update to [RFC 6053](#), which captured the results of the first ForCES interoperability test. The first test on ForCES was held in July 2008 at the University of Patras, Greece. That test focused on validating the basic semantics of the ForCES protocol and ForCES FE model.

1.1. ForCES Protocol

The ForCES protocol works in a master-slave mode in which Forwarding Elements (FEs) are slaves and Control Elements (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 Forwarding Element (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. [RFC 5811](#) specifies an SCTP-Based Transport Mapping Layer (TML) for ForCES protocol, which is a mandated TML for ForCES. See [RFC 5811](#) for more details.

1.4. Definitions

This document follows the terminology defined by ForCES related documents, including [RFC3654](#), [RFC3746](#), [RFC5810](#), [RFC5811](#), [RFC5812](#), [RFC5813](#), etc.

2. Overview

2.1. 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:

- o Zhejiang Gongshang University/Hangzhou BAUD Corporation of Information and Networks Technology (Hangzhou BAUD Networks), China. This implementation is referred to as "ZJSU" or in some cases "Z" in the document for the sake of brevity.
- o NTT Corporation, Japan. This implementation is referred to as "NTT" or in some cases "N" in the document for the sake of brevity.
- o The University of Patras, Greece. This implementation is referred to as "UoP" or in some cases "P" 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 5.1](#)).

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 would be implementing any missing features in the future.

[2.2.](#) Testbed Configuration

[2.2.1.](#) Participants Access

NTT and ZJSU 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](#)] .

2.2.2. Testbed Configuration

CEs and FEs from ZJSU and NTT 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 routing protocol testing machine, which acts as a router running OSPF and RIP and exchanged routing protocol messages with ForCES routers in the network. Connection to the Internet was via an ADSL channel.

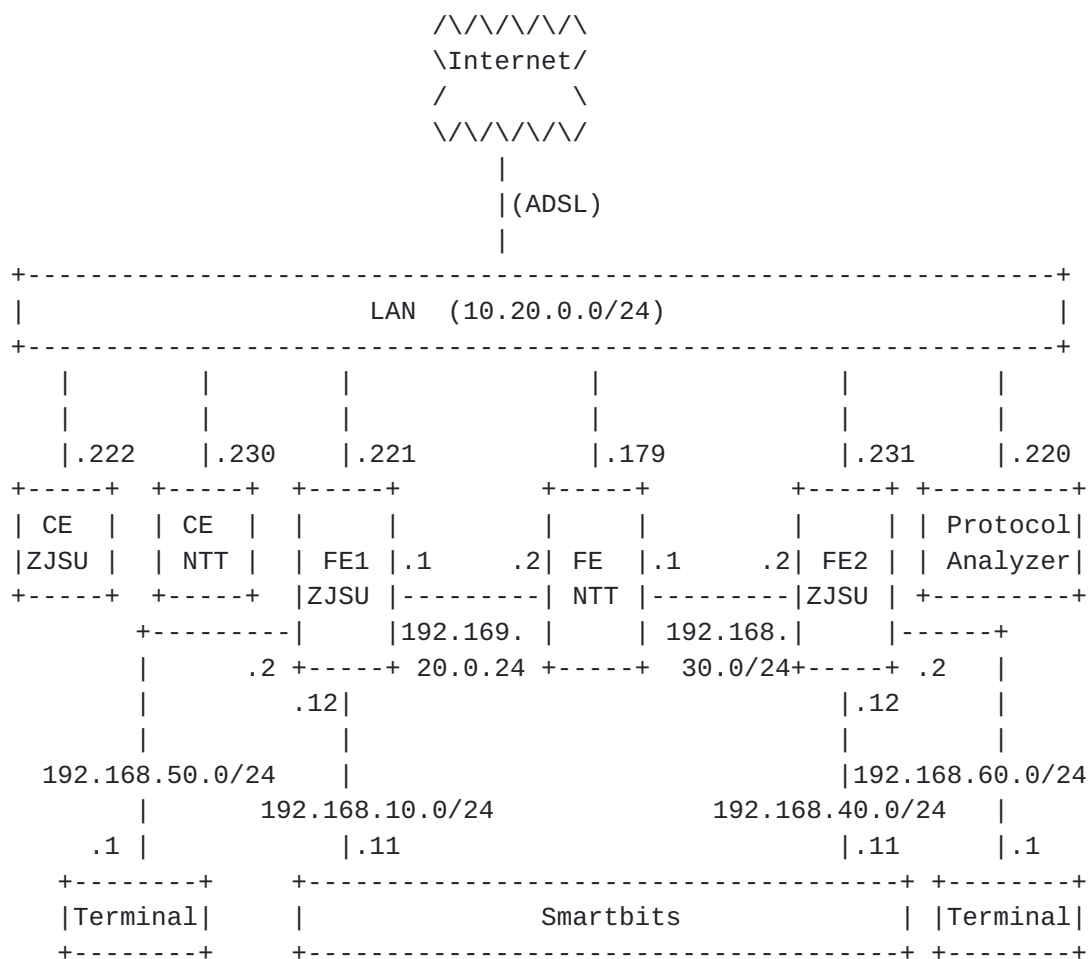


Figure 2: Testbed Configuration Located in ITL Lab, China

CE and FE from the UoP implementation were located within the University of Patras, Greece, and were connected together using LAN as shown in Figure 3. Connection to the Internet was via a VPN channel.




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

The testbeds above were then able to satisfy requirements of all interoperability test scenarios in this document.

3. Scenarios

3.1. Scenario 1 - LFB Operation

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

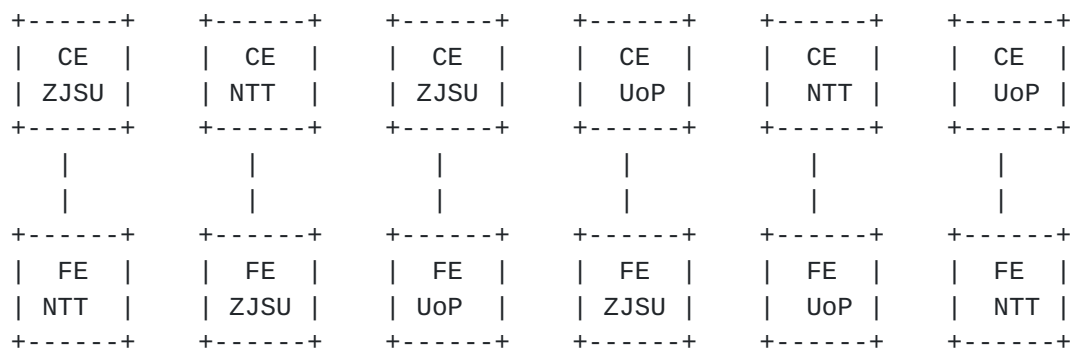


Figure 4: Scenario for LFB Operation

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

The test scenario was designed with the following purposes:

Firstly, the scenario was designed to verify all kinds of protocol messages with their complex data formats, which were defined in [RFC 5810](#). Specially, we tried 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 was designed to verify the definition of ForCES LFB Library [[I-D.ietf-forces-lfb-lib-03](#)], which defined a base set of ForCES LFB classes for typical router functions. Successful test under this scenario would help the validity of the LFB definitions.

3.2. Scenario 2 - TML with IPsec

This scenario was designed to implement a TML with IPsec, which was the requirement by [RFC 5811](#). TML with IPsec was not implemented and tested in the first ForCES interoperability test as reported by [RFC 6053](#). For this reason, in this 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 allowed us to verify whether all interactions between CE and FE could be made correctly under an IPsec TML environment.

The connection diagram for this scenario is shown as Figure 5. Because an unfortunate problem with the test system in UoP prevented the deployment of IPsec over TML, this test only took place between the test systems in ZJSU and NTT.




Figure 5: Scenario for LFB Operation with TML over IPsec

In this scenario, ForCES TML was run over the IPsec channel. Implementers joined in this interoperability used the same third-party software 'Racoon' [[Racoon](#)] to establish the IPsec channel.

The Racoon in NetBSD is an IKE daemon performing the IPsec Key Exchange (IKE) with the peers. Both IKE v1 and v2 are supported by Racoon in linux 2.6, and the v2 was adopted in the interop test. SAD and SPD were necessary for the test, setups of which were in the Racoon configuration file. ESP was specified in SAD and SPD in the Racoon configuration file.

ZJSU and NTT made a successful test with the scenario, and the following items were realized:

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

[3.3.](#) Scenario 3 - CE High Availability

CE High Availability (CEHA) was tested based on the ForCES CEHA document [[I-D.ietf-forces-ceha-00](#)]

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 were:

- 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.

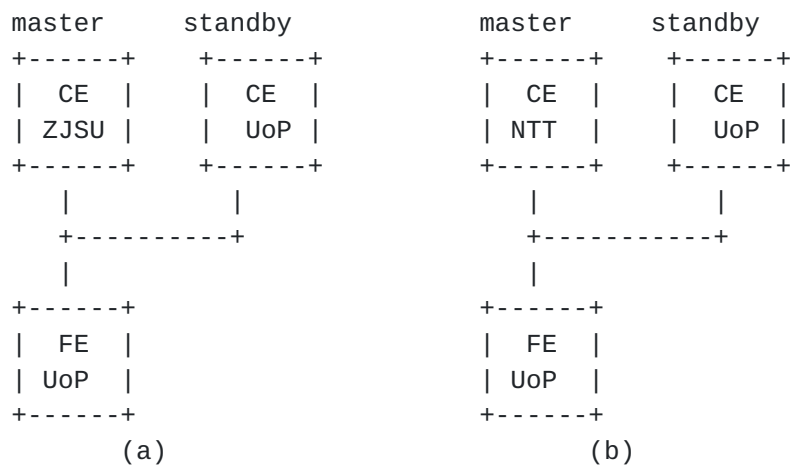


Figure 6: Scenario for CE High Availability

In this scenario one FE was connected and associated to a master CE and a backup CE. In the pre-association phase, the FE would be configured to have ZJSU's or NTT's CE as master CE and UoP's CE as standby CE. The CEFailoverPolicy component of the FE Protocol Object LFB that specified whether the FE was 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 was set to 2 or 3, the FE (in the post-association phase) would attempt to connect and associate with the standby CE.

When the master CE was deemed disconnected, either by a TearDown, Loss of Heartbeats or physically disconnected, the FE would assume that the standby CE was now the master CE. The FE would 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 would 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 considered itself as the master as well.

The steps of the CEHA test scenario were 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.

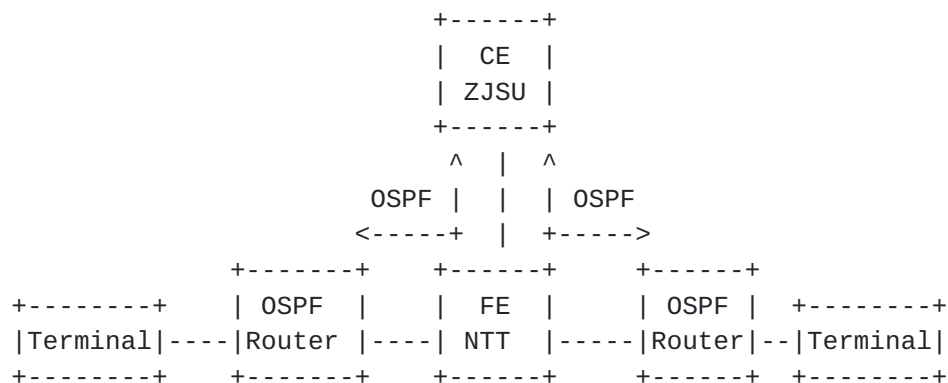
3.4. Scenario 4 - Packet forwarding

This test scenario was to verify LFBs like RedirectIn, RedirectOut, IPv4NextHop, IPv4UcastLPM defined by the ForCES LFB library document [[I-D.ietf-forces-lfb-lib-03](#)], 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)



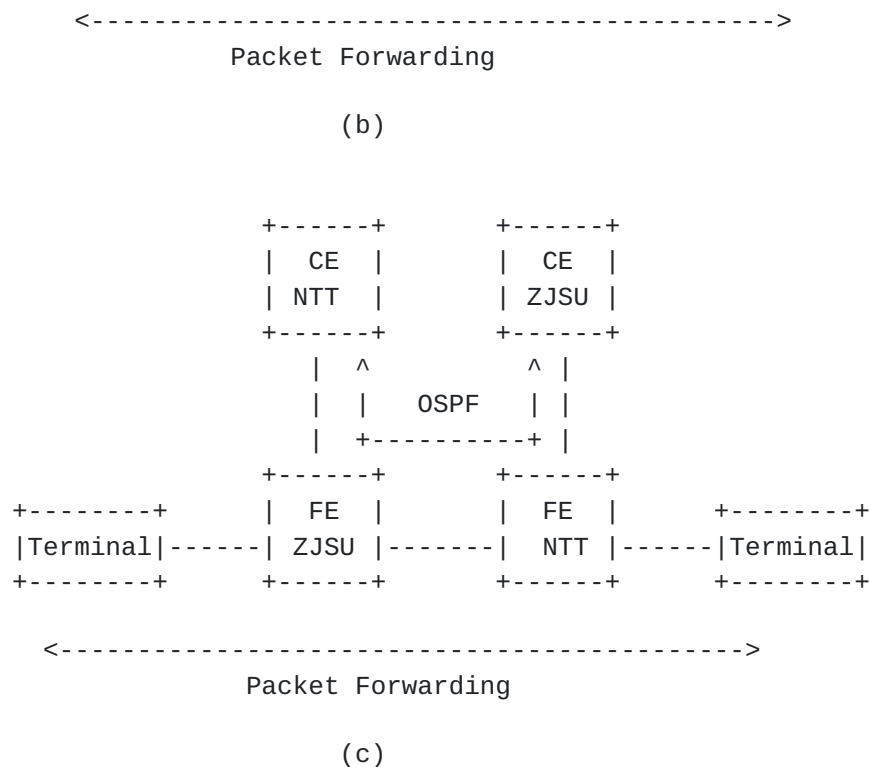


Figure 7: Scenario for IP Packet forwarding

In case (a), a CE by NTT was connected to an FE by ZJSU to form a ForCES router. A Smartbits test machine with its routing protocol software were used to simulate an OSPF router and were connected with the ForCES router to try to exchange OSPF hello packets and LSA packets among them. Terminals were 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 ZJSU was connected to an FE by NTT to form a ForCES router. Two routers running OSPF were simulated and connected to the ForCES router to test if the ForCES router could support OSPF protocol and support packet forwarding.

In case (c), two ForCES routers were constructed. One was with CE by NTT and FE by ZJSU and the other was opposite. OSPF and packet forwarding were 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.

4. Test Results

4.1. LFB Operation Test

The test result is as reported by Figure 8. For the convenience sake, as mentioned earlier, abbreviations of 'Z' in the table means implementation from ZJSU, 'N' implementation from NTT, and 'P' implementation from UoP.

Test#	CE	FE(s)	Oper	LFB	Component /Capability	Result
1	Z	N	GET	FEObject	LFBTopology	Success
	N	Z				Success
	P	Z				Success
	N	P				Success
	P	N				Success
2	Z	N	GET	FEObject	LFBSelector	Success
	N	Z				Success
	Z	P				Success
	P	Z				Success
	N	P				Success
	P	N				Success
3	Z	N	GET	EtherPHYCop	PHYPortID	Success
	N	Z				Success
	Z	P				Success
	P	Z				Success
	N	P				Success
	P	N				Success
4	Z	N	GET	EtherPHYCop	AdminStatus	Success
	N	Z				Success
	Z	P				Success
	P	Z				Success
	N	P				Success
	P	N				Success
5	Z	N	GET	EtherPHYCop	OperStatus	Success

		N	Z				Success
		Z	P				Success
		P	Z				Success
		N	P				Success
		P	N				Success
6		Z	N	GET	EtherPHYCop	AdminLinkSpeed	Success
		N	Z				Success
		Z	P				Success
		P	Z				Success
		N	P				Success
		P	N				Success
7		Z	N	GET	EtherPHYCop	OperLinkSpeed	Success
		N	Z				Success
		Z	P				Success
		P	Z				Success
		N	P				Success
		P	N				Success
8		Z	N	GET	EtherPHYCop	AdminDuplexSpeed	Success
		N	Z				Success
		Z	P				Success
		P	Z				Success
		N	P				Success
		P	N				Success
9		Z	N	GET	EtherPHYCop	OperDuplexSpeed	Success
		N	Z				Success
		Z	P				Success
		P	Z				Success
		N	P				Success
		P	N				Success
10		Z	N	GET	EtherPHYCop	CarrierStatus	Success
		N	Z				Success
		Z	P				Success
		P	Z				Success
		N	P				Success
		P	N				Success
11		Z	N	GET	EtherMACIn	AdminStatus	Success
		N	Z				Success
		Z	P				Success
		P	Z				Success
		N	P				Success
		P	N				Success

12	Z	N	GET	EtherMACIn	LocalMacAddresses	Success
	N	Z				Success
	Z	P				Success
	P	Z				Success
	N	P				Success
	P	N				Success
13	Z	N	GET	EtherMACIn	L2Bridging PathEnable	Success
	N	Z				Success
	Z	P				Success
	P	Z				Success
	N	P				Success
	P	N				Success
14	Z	N	GET	EtherMACIn	PromiscuousMode	Success
	N	Z				Success
	Z	P				Success
	P	Z				Success
	N	P				Success
	P	N				Success
15	Z	N	GET	EtherMACIn	TxFlowControl	Success
	N	Z				Success
	Z	P				Success
	P	Z				Success
	N	P				Success
	P	N				Success
16	Z	N	GET	EtherMACIn	RxFlowControl	Success
	N	Z				Success
	Z	P				Success
	P	Z				Success
	N	P				Success
	P	N				Success
17	Z	N	GET	EtherMACIn	MACInStats	Success
	N	Z				Success
	Z	P				Success
	P	Z				Success
	N	P				Success
	P	N				Success
18	Z	N	GET	EtherMACOut	AdminStatus	Success
	N	Z				Success
	Z	P				Success
	P	Z				Success
	N	P				Success
	P	N				Success

19	Z	N	GET	EtherMACOut	MTU	Success
	N	Z				Success
	Z	P				Success
	P	Z				Success
	N	P				Success
	P	N				Success
20	Z	N	GET	EtherMACOut	TxFlowControl	Success
	N	Z				Success
	Z	P				Success
	P	Z				Success
	N	P				Success
	P	N				Success
21	Z	N	GET	EtherMACOut	TxFlowControl	Success
	N	Z				Success
	Z	P				Success
	P	Z				Success
	N	P				Success
	P	N				Success
22	Z	N	GET	EtherMACOut	MACOutStats	Success
	N	Z				Success
	Z	P				Success
	P	Z				Success
	N	P				Success
	P	N				Success
23	Z	N	GET	ARP	PortV4AddrInfoTable	Success
	N	Z				Success
	Z	P				Success
	P	Z				Success
	N	P				Success
	P	N				Success
24	Z	N	SET	ARP	PortV4AddrInfoTable	Success
	N	Z				Success
	Z	P				Success
	P	Z				Success
	N	P				Success
	P	N				Success
25	Z	N	DEL	ARP	PortV4AddrInfoTable	Success
	N	Z				Success
	Z	P				Success
	P	Z				Success
	N	P				Success

		P	N				Success
26	Z	N	SET	EtherMACIn	LocalMACAddresses	Success	
	N	Z				Success	
	Z	P				Success	
	P	Z				Success	
	N	P				Success	
	P	N				Success	
27	Z	N	SET	EtherMACIn	MTU	Success	
	N	Z				Success	
	Z	P				Success	
	P	Z				Success	
	N	P				Success	
	P	N				Success	
28	Z	N	SET	IPv4NextHop	IPv4NextHopTable	Success	
	N	Z				Success	
	Z	P				Success	
	P	Z				Success	
	N	P				Success	
	P	N				Success	
29	Z	N	SET	IPv4UcastLPM	IPv4PrefixTable	Success	
	N	Z				Success	
	Z	P				Success	
	P	Z				Success	
	N	P				Success	
	P	N				Success	
30	Z	N	DEL	IPv4NextHop	IPv4NextHopTable	Success	
	N	Z				Success	
	Z	P				Success	
	P	Z				Success	
	N	P				Success	
	P	N				Success	
31	Z	N	DEL	IPv4UcastLPM	IPv4PrefixTable	Success	
	N	Z				Success	
	Z	P				Success	
	P	Z				Success	
	N	P				Success	
	P	N				Success	
32	Z	N	SET	EtherPHYCop	AdminStatus	Success	
	N	Z				Success	
	Z	P				Success	
	P	Z				Success	

	N	P				Success
	P	N				Success
33	Z	N	SET	Ether Classifier	VlanInputTable	Success
	N	Z				Success
	Z	P				Success
	P	Z				Success
	N	P				Success
	P	N				Success
34	Z	N	DEL	Ether Classifier	VlanInputTable	Success
	N	Z				Success
	Z	P				Success
	P	Z				Success
	N	P				Success
	P	N				Success
35	Z	N	SET	Ether Encapsulator	VlanOutputTable	Success
	N	Z				Success
	Z	P				Success
	P	Z				Success
	N	P				Success
	P	N				Success
36	Z	N	DEL	Ether Encapsulator	VlanOutputTable	Success
	N	Z				Success
	Z	P				Success
	P	Z				Success
	N	P				Success
	P	N				Success

Figure 8: LFB Operation Test Results

Note on test #1 and #2:

On the wire format of encapsulation on array, only the case of FULLDATA vs SPARSEDATA was tested.

It is very common for CE to get information of FEobject LFB in FE so as to know status on all active LFBs in the FE. Hence, the two tests were specifically designed.

4.2. TML with IPsec Test

In this scenario, the ForCES TML was run over IPsec. Implementers joined this interoperability test used the same third-party tool software 'Racoon' [[Racoon](#)] to establish IPsec channel. Typical LFB operation tests as in Scenario 1 were repeated with the IPsec enabled TML.

As mentioned, this scenario only took place between implementers of ZJSU and NTT.

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

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

Figure 9: TML with IPsec Test Results

4.3. CE High Availability Test

In this scenario, one FE connected and associated with a master CE and a backup CE. When the master CE was deemed disconnected, the FE would attempt to find another associated CE to become the master CE.

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

Due to a bug in one of the FEs, an 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 was disconnected.

While the bug was fixed and the CEHA scenario was completed successfully, the authors felt 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.

4.4. Packet Forwarding Test

As described in the ForCES LFB library [[I-D.ietf-forces-lfb-lib-03](#)], 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 were used. RedirectIn and RedirectOut LFBs redirected OSPF hello and LSA packets from and to CE. A Smartbits test machine was used to simulate an OSPF router and exchange the OSPF hello and LSA packets with CE in ForCES router.

In Figure 7, case (a) and case (b) both need a RedirectIn LFB to send OSPF packets generated by CE to FE by use of ForCES packet redirect messages. The OSPF packets were further sent to an outside OSPF Router by the FE via forwarding LFBs including IPv4NextHop and IPv4UcastLPM LFBs. A RedirectOut LFB in the FE was 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 could generate new routes and load them to routing table in FE. The FE was then able to forward packets according to the routing table.

The test results are as shown in Figure 10

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

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


```

|         |         |         |         |         | IPv4NextHop |
|         |         |         |         |         | IPv4UcastLPM|
+-----+-----+-----+-----+-----+-----+

```

Figure 10: Packet Forwarding Test Results

Note on test #1 and #2:

The implementer found a multicast route pointing to localhost had to be manually set before a redirect channel could work normally.

Note on test #3 to #9:

During the test, OSPF packets received from CE were found by Ethereal/Wireshark with checksum errors in FE. Because the test time was quite limited, implementer of the CE did not try to make efforts to find and solve the checksum error, instead, the FE had tried to correct the checksum in order not to let the Smartbits drop the packets. Note that such solution does not affect the test results.

Comment on Test #16 and #17:

The two test items failed. Note that Test #7 and #8 were identical to the two tests, only with CE and FE implementers were exchanged. Moreover, test #12 and #13 showed that the redirect channel worked well. Therefore, it can be reasonably inferred that the problem caused the failure was from the implementations, rather than from the ForCES protocol itself or from misunderstanding of implementations on the protocol specification. Although the failure made the OSPF interoperability test incomplete, it did not show interoperability problem. More test work is needed to verify the OSPF interoperability.

5. Discussions

5.1. On Data Encapsulation Format

In the first day of the test, it was found that the LFB inter-operations about tables all failed. It was eventually found the failure was because that different data encapsulation methods for ForCES protocol messages were taken by different implementations. The issue is described in detail 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 testers 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) should 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 has shown 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.

6. 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@zjsu.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@zjsu.edu.cn

7. Acknowledgements

The authors thank the following test participants:

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

The authors also thank very much to Adrian Farrel, Joel Halpern, Ben Campbell, Nevil Brownlee, and Sean Turner for their important help in the document publication process.

8. IANA Considerations

This memo includes no request to IANA.

9. 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. Test results of TML with IPsec supported have been shown in [Section 4.2](#) in this document.

The tests described in this document used only simple password security mode. Testing using more sophisticated security is for future study.

Further testing using key agility is encouraged. The tests reported here used SCTP TML running over an IPsec tunnel which was established by Racoon. Key negotiation formed part of this process, but we believe that the SCTP TML used does not include key agility or renegotiation.

10. References

10.1. 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.

10.2. 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-00]

Ogawa, K., Wang, W., Haleplidis, E., and J. Salim, "ForCES Intra-NE High Availability", [draft-ietf-forces-ceha-00](#) (work in progress) [RFC Editor Note: This reference is intended to indicate a specific version of an Internet-Draft that was used during interop testing. Please Do NOT update this reference to a more recent version of the draft or to an RFC. Please remove this note before publication] , October 2010.

[I-D.ietf-forces-lfb-lib-03]

Wang, W., Haleplidis, E., Ogawa, K., Li, C., and J. Halpern, "ForCES Logical Function Block (LFB) Library", [draft-ietf-forces-lfb-lib-03](#) (work in progress) [RFC Editor Note: This reference is intended to indicate a specific version of an Internet-Draft that was used during interop testing. Please Do NOT update this reference to a more recent version of the draft or to an RFC. Please remove this note before publication] , December 2010.

[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.

[Racoon] , "Racoon in NetBSD is a well-known IKE daemon performing the IPsec Key Exchange (IKE) with the peers", <http://www.netbsd.org/docs/network/ipsec/rasvpn.html> , .

[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@zjsu.edu.cn

Kentaro Ogawa
NTT Corporation
Tokyo
Japan

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

Evangelos Haleplidis
University of Patras
Department of Electrical & Computer Engineering
Patras 26500
Greece

Email: ehalep@ece.upatras.gr

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

Email: gmyyqno1@zjsu.edu.cn

Jamal Hadi Salim
Mojatatu Networks
Ottawa
Canada

Email: hadi@mojatatu.com

