

| | | |
|---------------------------------|-------------------------------|--|
| Internet Engineering Task Force | E. Haleplidis | |
| Internet-Draft | University of Patras | |
| Intended status: Informational | K. Ogawa | |
| Expires: January 1, 2010 | NTT Corporation | |
| | X. Wang | |
| | Huawei Technologies Co., Ltd. | |
| | C. Li | |
| | Zhejiang Gongshang University | |
| | June 30, 2009 | |

[TOC](#)

ForCES Interoperability Draft draft-ietf-forces-interoperability-02

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

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/1id-abstracts.txt>.

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

This Internet-Draft will expire on January 1, 2010.

Copyright Notice

Copyright (c) 2009 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 in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Abstract

This document describes the details of the interoperability test of the Forward and Control Element Separation (ForCES) protocol that will take place in the University of Patras in Rio, Greece, in the third week of July 2009. This informational draft provides necessary information, for all parties who wish to participate in the interoperability test.

Table of Contents

| | |
|-----------------------|---|
| 1. | Terminology and Conventions |
| 1.1. | Requirements Language |
| 2. | Introduction |
| 2.1. | ForCES Protocol |
| 2.2. | ForCES Model |
| 2.3. | Transport mapping layer |
| 3. | Definitions |
| 4. | Date, Location and Access |
| 4.1. | Date |
| 4.2. | Location |
| 4.3. | Access |
| 5. | Testbed architecture |
| 5.1. | Local configuration |
| 5.2. | Distributed configuration |
| 6. | Scenarios |
| 6.1. | Scenario 1 - Pre-association Setup |
| 6.2. | Scenario 2 - TML priority channels connection |
| 6.3. | Scenario 3 - Association Setup - Association Complete |
| 6.4. | Scenario 4 - CE query |
| 6.5. | Scenario 5 - Heartbeat monitoring |
| 6.6. | Scenario 6 - Simple Config Command |
| 6.7. | Scenario 7 - Association Teardown |
| 7. | Acknowledgements |
| 8. | IANA Considerations |
| 9. | Security Considerations |
| 10. | References |
| 10.1. | Normative References |
| 10.2. | Informative References |
| § | Authors' Addresses |

1.1. Requirements Language

[TOC](#)

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 [\[RFC2119\] \(Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," March 1997.\)](#).

2. Introduction

[TOC](#)

Forwarding and Control Element Separation (ForCES) defines an architectural framework and associated protocols to standardize information exchange between the control plane and the forwarding plane in a ForCES Network Element (ForCES NE). [\[RFC3654\] \(Khosravi, H. and T. Anderson, "Requirements for Separation of IP Control and Forwarding," November 2003.\)](#) has defined the ForCES requirements, and [\[RFC3746\] \(Yang, L., Dantu, R., Anderson, T., and R. Gopal, "Forwarding and Control Element Separation \(ForCES\) Framework," April 2004.\)](#) has defined the ForCES framework.

2.1. ForCES Protocol

[TOC](#)

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 [FE-protocol \(Dong, L., Doria, A., Gopal, R., HAAS, R., Salim, J., Khosravi, H., and W. Wang, "ForCES Protocol Specification," March 2009.\)](#) [I-D.ietf-forces-protocol] for further information.

2.2. ForCES Model

[TOC](#)

The [FE-MODEL \(Halpern, J. and J. Salim, "ForCES Forwarding Element Model," October 2008.\)](#) [I-D.ietf-forces-model] 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.

2.3. Transport mapping layer

[TOC](#)

The TML transports the 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 [\[RFC3654\] \(Khosravi, H. and T. Anderson, "Requirements for Separation of IP Control and Forwarding," November 2003.\)](#) will be used.

3. Definitions

[TOC](#)

This document follows the terminology defined by the ForCES Requirements in [\[RFC3654\] \(Khosravi, H. and T. Anderson, "Requirements for Separation of IP Control and Forwarding," November 2003.\)](#) and by the ForCES framework in [\[RFC3746\] \(Yang, L., Dantu, R., Anderson, T., and R. Gopal, "Forwarding and Control Element Separation \(ForCES\) Framework," April 2004.\)](#). The definitions below 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.

CE Manager (CEM) - A logical entity responsible for generic CE management tasks. It is particularly used during the pre-association phase to determine with which FE(s) a CE should communicate. This process is called FE discovery and may involve the CE manager learning the capabilities of available FEs.

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.

FE Manager (FEM) - A logical entity responsible for generic FE management tasks. It is used during pre-association phase to determine with which CE(s) an FE should communicate. This process is called CE discovery and may involve the FE manager learning

the capabilities of available CEs. An FE manager may use anything from a static configuration to a pre-association phase protocol (see below) to determine which CE(s) to use. Being a logical entity, an FE manager might be physically combined with any of the other logical entities such as FEs.

ForCES Network Element (NE) - An entity composed of one or more CEs and one or more FEs. To entities outside an NE, the NE represents a single point of management. Similarly, an NE usually hides its internal organization from external entities.

LFB (Logical Function 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 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.

FE Topology - A representation of how the multiple FEs within a single NE are interconnected. Sometimes this is called inter-FE topology, to be distinguished from intra-FE topology (i.e., LFB topology).

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 Component - 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 (see below).

LFB Topology - Representation of how the LFB instances are logically interconnected and placed along the datapath within one

FE. Sometimes it is also called intra-FE topology, to be distinguished from inter-FE topology.

Pre-association Phase - The period of time during which an FE Manager and a CE Manager are determining which FE(s) and CE(s) should be part of the same network element.

Post-association Phase - The period of time during which an FE knows which CE is to control it and vice versa. This includes the time during which the CE and FE are establishing communication with one another.

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. This document defines the specifications for this ForCES protocol.

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.

4. Date, Location and Access

[TOC](#)

4.1. Date

[TOC](#)

The date that the Interoperability draft will take place has been specified at 15-16/07/2009, one and a half week before IETF 75, in Stockholm.

[TOC](#)

4.2. Location

Patras is a major harbor of Greece connecting it with Italy.
The University of Patras is located in Rio, 10km east out of Patras.
The following coordinates mark the Electrical Engineering building in the University.

*North: 38°17'15.99"

*East: 21°47'19.28"

4.3. Access

[TOC](#)

The best way to come to Greece is by plane to the Athens International Airport.

From there there are three ways to arrive in the University of Patras.

1. Renting a car and driving to the University. It is a maximum 2:30 hours drive from the airport.
2. Via coach station. Get from the airport to the coach station via X93 bus towards the Kifissos Coach Station. At the Coach Station there are buses to Patras every 30 minutes. The Bus to Patras may take about 2:30 - 3:00 hours, and the ride of the X93 bus may take about 30 mins - 1hour depending on the traffic, so it's about 3:30 - 4:30 hours away with the wait at the Coach Station.
3. Via Train. It is recommended you already have booked your ticket beforehand as there are not many trains going to Patras, and mostly are booked in advanced. It is not recommended that you take the train to Patras, as you have to change at least 2 trains. In order to reach Patras from the Athens International Airport you need to take the Suburban Rail to Neratziotissa. From there you must take ISAP to Pireaus. There you must change again to Suburban Rail to reach Kiato. From Kiato you can catch a train to Patras. It will take you at least 5 hours to reach Patras.

[TOC](#)

5. Testbed architecture

Most FEs and CEs should be located locally at the University of Patras premises. But if some parties would like to participate but cannot attend the interoperability test locally a connection over the internet MAY be created.

The actual test will take place between FEs and CEs of different implementors with different permutations.

All protocol messages of each scenario will be monitored using a protocol network analyzer to test validity. Two tools shall be used:

*A modified tcpdump [\[tcpdump\]](#) (, ["Tcpdump is a linux protocol analyzer. The specific tcpdump that will be used is a modified tcpdump, by Jamal Hadi Salim, that can analyze and decode the ForCES protocol messages.," .\)](#)).

*A modified Ethereal [\[ethereal\]](#) (, ["Ethereal is a protocol analyzer. The specific ethereal that will be used is an updated Ethereal, by Fenggen Jia, that can analyze and decode the ForCES protocol messages.," .\)](#)).


All NE's in all the scenarios will be comprised of one CE and one FE from different implementors.

5.1. Local configuration


[TOC](#)

Hardware/Software (CEs and FEs) that will be located within the University of Patras premises, will be connected together using switches.

The scenarios will be tested with only one CE associated with one or multiple FEs from different implementors. The CE and the FE(s) will be connected in one LAN as shown in the following figure.



All scenarios will be tested more than once with permutation of the CE from different implementors. In the next permutation, the setup will be as shown in the following figure.



5.2. Distributed configuration

[TOC](#)

For parties that cannot participate, public IPs can be provided and associations can be achieved over the internet as seen in the following figure.

```

+-----+ +-----+ /\ /\ /\ /\ +-----+ +-----+
|FE/CE| |Implementor| \Internet/ |University| |FE/CE|
|ImplX|---| Router |---/ \---| Router |---|ImplY|
+-----+ +-----+ /\ /\ /\ /\ +-----+ +-----+

```

For interoperability issues, all CEs and FEs MUST implement no security even in the TML. For security, firewalls MUST be used that will allow only the specific IPs and the SCTP ports defined in the [SCTP-TML draft \(Salim, J. and K. Ogawa, "SCTP based TML \(Transport Mapping Layer\) for ForCES protocol," January 2009.\)](#) [I-D.ietf-forces-sctptml].

6. Scenarios

[TOC](#)

Since the main goal of this interoperability test is to test the basic protocol functionality, we will limit the test parameters. Therefore:

1. In the Association Setup Message, all report messages will be ignored.
 2. In the Association Setup Phase, the messages, FEO OperEnable Event (FE to CE), Config FEO Adminup (CE to FE) and FEO Config-Resp (FE to CE) will be ignored. The CE will assume that the FE is enabled once the LFBSelectors has been queried.
 3. Only FullDataTLVs are going to be used and not SparseData TLV's.
 4. There will be no transaction operations.
 5. Each message shall have only one LFBSelector TLV, one Operation TLV and one PathDataTLV per message when these are used.
-

6.1. Scenario 1 - Pre-association Setup

[TOC](#)

While the Pre-association setup is not in the ForCES current scope it is an essential step before CEs and FEs communicate. As the first part in a successful CE-FE connection the participating CEs and FEs should be able to be configured.

In the Pre-association Phase the following configuration items MUST be setup regarding the CEs:

*The CE ID.

- *The FE IDs that will be connected to this CE

- *The IP of the FEs that will connect

- *The TML priority ports.

In the Pre-association Phase the following configuration items MUST be setup regarding the FEs:

- *The FE ID.

- *The CE ID that this FE will be connecting to.

- *The IP of the CE that will connect to

- *The TML priority ports.

Once each element is setup and configured, Scenario 1 is successful.

6.2. Scenario 2 - TML priority channels connection

[TOC](#)

For the current interoperability test, the SCTP will be used as TML. The TML connection with the associating element is needed for the scenario 2 to be successful.

The [SCTP-TML draft \(Salim, J. and K. Ogawa, "SCTP based TML \(Transport Mapping Layer\) for ForCES protocol," January 2009.\)](#)

[I-D.ietf-forces-sctptml] defines 3 priority channels, with specific ports:

- *High priority - Port number: 6700

- *Medium priority - Port number: 6701

- *Lower priority - Port number: 6702

Once these channels have been established with each associated element, will the Scenario 2 be successful.

6.3. Scenario 3 - Association Setup - Association Complete

[TOC](#)

Once the Pre-association phase has been complete in the previous 2 scenarios, CEs and FEs are ready to communicate using the ForCES protocol, and enter the Association Setup stage. In this stage the FEs attempts to join the NE. The following ForCES protocol messages will be exchanged for each CE-FE pair in the specified order:

*Association Setup Message (from FE to CE)

*Association Setup Response Message (from CE to FE)

*Query Message: FE0 LFBSelectors(from CE to FE)

*Query Response: FE0 LFBSelectors response (from FE to CE)

Once the associations has been initialized scenario 3 will have been successful.

6.4. Scenario 4 - CE query

[TOC](#)

Once the Association Phase stage has been complete, the FEs and CEs will enter the Established stage. In this stage the FE is continuously updated or queried. The CE should query the FE a specific value from the FE Object LFB and from the FE Protocol LFB. An example from the FE Protocol LFB is the HeartBeat Timer (FEHI) and from the FE Object LFB is the State of the LFB (FEState)

The following ForCES protocol messages will be exchanged:

*Query Message

*Query Response Message

6.5. Scenario 5 - Heartbeat monitoring

[TOC](#)

The Heartbeat (HB) Message is used for one ForCES element (FE or CE) to asynchronously notify one or more other ForCES elements in the same ForCES NE on its liveness. The default configuration of the Heartbeat Policy of the FE is set to 0 which means, that the FE should not generate any Heartbeat messages. the CE is responsible for checking FE liveness by setting the PL header ACK flag of the message it sends to AlwaysACK. In this Scenario the CE should send a Heartbeat message with the ACK flag set to AlwaysACK and the FE should respond.

The following ForCES protocol messages will be exchanged:

*Heartbeat Message

[TOC](#)

6.6. Scenario 6 - Simple Config Command

A config message is sent by the CE to the FE to configure LFB components in the FE. A simple config command easily visible and metered would be to change the Heartbeat configuration. This will be done in two steps:

1. Change the FE Heartbeat Policy (FEHBPoly) to value 1, to force the FE to send heartbeats.
2. After some heartbeats from the FE, the FE Heartbeat Interval (FEHI) will be changed.

The following ForCES protocol messages will be exchanged:

*Config Message

*Config Response Message

6.7. Scenario 7 - Association Teardown

[TOC](#)

In the end, the association must be terminated. There are two scenarios by which the association maybe terminated:

1. Normal tear down by exchanging Association Teardown Message
2. Irregular tear down by stopping heartbeats from a FE or a CE.
3. Irregular tear down by externally shutting down/rebooting a FE or a CE.

All scenarios may be tested in the interoperability test.

The following ForCES protocol messages will be exchanged:

*Association Teardown Message

7. Acknowledgements

[TOC](#)

The authors of this draft would like to acknowledge and thank the chair of the ForCES working group Jamal Hadi Salim.

8. IANA Considerations

[TOC](#)

This memo includes no request to IANA.

9. Security Considerations

[TOC](#)

Section 9 of the [FE-protocol](#) (Dong, L., Doria, A., Gopal, R., HAAS, R., Salim, J., Khosravi, H., and W. Wang, "ForCES Protocol Specification," March 2009.) [I-D.ietf-forces-protocol] specifies security considerations of the ForCES protocol. For this interoperability test, no security MUST be chosen even for the distributed architecture.

10. References

[TOC](#)

10.1. Normative References

[TOC](#)

| | |
|----------------------------|--|
| [I-D.ietf-forces-model] | Halpern, J. and J. Salim, " ForCES Forwarding Element Model ," draft-ietf-forces-model-16 (work in progress), October 2008 (TXT). |
| [I-D.ietf-forces-protocol] | Dong, L., Doria, A., Gopal, R., HAAS, R., Salim, J., Khosravi, H., and W. Wang, " ForCES Protocol Specification ," draft-ietf-forces-protocol-22 (work in progress), March 2009 (TXT). |
| [I-D.ietf-forces-sctptml] | Salim, J. and K. Ogawa, " SCTP based TML (Transport Mapping Layer) for ForCES protocol ," draft-ietf-forces-sctptml-02 (work in progress), January 2009 (TXT). |

10.2. Informative References

[TOC](#)

| | |
|-----------|--|
| [RFC2119] | Bradner, S. , " Key words for use in RFCs to Indicate Requirement Levels ," BCP 14, RFC 2119, March 1997 (TXT , HTML , XML). |
| [RFC2629] | Rose, M. , " Writing I-Ds and RFCs using XML ," RFC 2629, June 1999 (TXT , HTML , XML). |
| [RFC3552] | Rescorla, E. and B. Korver , " Guidelines for Writing RFC Text on Security Considerations ," BCP 72, RFC 3552, July 2003 (TXT). |

| | |
|------------|---|
| [RFC3654] | Khosravi, H. and T. Anderson, " Requirements for Separation of IP Control and Forwarding ," RFC 3654, November 2003 (TXT). |
| [RFC3746] | Yang, L., Dantu, R., Anderson, T., and R. Gopal, " Forwarding and Control Element Separation (ForCES) Framework ," RFC 3746, April 2004 (TXT). |
| [RFC5226] | Narten, T. and H. Alvestrand, " Guidelines for Writing an IANA Considerations Section in RFCs ," BCP 26, RFC 5226, May 2008 (TXT). |
| [ethereal] | " Ethereal is a protocol analyzer. The specific ethereal that will be used is an updated Ethereal, by Fenggen Jia, that can analyze and decode the ForCES protocol messages.. " |
| [tcpdump] | " Tcpdump is a linux protocol analyzer. The specific tcpdump that will be used is a modified tcpdump, by Jamal Hadi Salim, that can analyze and decode the ForCES protocol messages.. " |

Authors' Addresses

[TOC](#)

| | |
|--------|--|
| | Evangelos Haleplidis |
| | University of Patras |
| | Patras, |
| | Greece |
| Email: | ehalep@ece.upatras.gr |
| | |
| | Kentaro Ogawa |
| | NTT Corporation |
| | Tokyo, |
| | Japan |
| Email: | ogawa.kentaro@lab.ntt.co.jp |
| | |
| | Xin-ping Wang |
| | Huawei Technologies Co., Ltd. |
| | China |
| Email: | carly.wang@huawei.com |
| | |
| | Chuanhuang Li |
| | Zhejiang Gongshang University |
| | 18, Xuezheng Str., Xiasha University Town |
| | Hangzhou, 310018 |
| | P.R.China |
| Phone: | +86-571-28877751 |
| Email: | chuanhuang_li@pop.zjgsu.edu.cn |