

Internet Engineering Task Force (IETF)
Request for Comments: 6041
Category: Informational
ISSN: 2070-1721

A. Crouch
H. Khosravi
Intel
A. Doria, Ed.
LTU
X. Wang
Huawei
K. Ogawa
NTT Corporation
October 2010

Forwarding and Control Element Separation (ForCES)
Applicability Statement

Abstract

The Forwarding and Control Element Separation (ForCES) protocol defines a standard framework and mechanism for the interconnection between control elements and forwarding elements in IP routers and similar devices. In this document we describe the applicability of the ForCES model and protocol. We provide example deployment scenarios and functionality, as well as document applications that would be inappropriate for ForCES.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6041>.

Copyright Notice

Copyright (c) 2010 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.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1. Introduction	3
2. Purpose	4
3. Terminology	4
4. Applicability to IP Networks	4
4.1. Applicable Services	5
4.1.1. Association, Capability Discovery, and Information Exchange	5
4.1.2. Topology Information Exchange	6
4.1.3. Configuration	6
4.1.4. Routing Exchange	6
4.1.5. QoS Capabilities Exchange and Configuration	7
4.1.6. Security Exchange	7
4.1.7. Filtering Exchange and Firewalls	7
4.1.8. Encapsulation/Tunneling Exchange	7
4.1.9. NAT and Application-Level Gateways	7
4.1.10. Measurement and Accounting	7
4.1.11. Diagnostics	8
4.1.12. Redundancy and Failover	8
4.2. CE-FE Link Capability	8
4.3. CE/FE Locality	8
5. Security Considerations	9
6. ForCES Manageability	9
6.1. The NE as an Atomic Element	10
6.2. The NE as Composed of Manageable Elements	10
6.3. ForCES Protocol MIB	10
6.3.1. MIB Management of an FE	11
6.4. The FEM and CEM	12
7. Contributors	12
8. Acknowledgments	12
9. References	12
9.1. Normative References	12
9.2. Informative References	13

1. Introduction

The Forwarding and Control Element Separation (ForCES) protocol defines a standard framework and mechanism for the exchange of information between the logically separate functionality of the control and data forwarding planes of IP routers and similar devices. It focuses on the communication necessary for separation of control plane functionality such as routing protocols, signaling protocols, and admission control from data forwarding plane per-packet activities such as packet forwarding, queuing, and header editing.

This document defines the applicability of the ForCES mechanisms. It describes types of configurations and settings where ForCES is most appropriately applied. This document also describes scenarios and configurations where ForCES would not be appropriate for use.

2. Purpose

The purpose of the ForCES Applicability Statement is to capture the intent of the ForCES protocol [RFC5810] designers as to how the protocol could be used in conjunction with the ForCES model [RFC5812] and a Transport Mapping Layer [RFC5811].

3. Terminology

A set of concepts associated with ForCES was introduced in "Requirements for Separation of IP Control and Forwarding" [RFC3654] and in "Forwarding and Control Element Separation (ForCES) Framework" [RFC3746]. The terminology associated with these concepts and with the protocol elements in ForCES is defined in the "Forwarding and Control Element Separation (ForCES) Protocol Specification" [RFC5810].

The reader is directed to these documents for the conceptual introduction and for definitions, including the following acronyms:

- o CE: control element
- o CEM: CE Manager
- o FE: forwarding element
- o FEM: FE Manager
- o ForCES: Forwarding and Control Element Separation protocol
- o LFB: Logical Function Block
- o NE: ForCES network element
- o TML: Transport Mapping Layer

4. Applicability to IP Networks

This section lists the areas of ForCES applicability in IP network devices. Some relatively low-end routing systems may be implemented on simple hardware that performs both control and packet forwarding functionality. ForCES may not be useful for such devices.

Higher-end routing systems typically distribute work amongst several interface-processing elements, and these devices (FEs) therefore need to communicate with the control element(s) to perform their job. A higher-end router may also distribute control processing amongst several processing elements (CEs). ForCES provides a standard way to do this communication. ForCES also provides support for high-availability configurations that include a primary CE and one or more secondary CEs.

The remainder of this section lists the applicable services that ForCES may support, applicable FE functionality, applicable CE-FE link scenarios, and applicable topologies in which ForCES may be deployed.

4.1. Applicable Services

In this section we describe the applicability of ForCES for the following control-forwarding-plane services:

- o Association, Capability Discovery, and Information Exchange
- o Topology Information Exchange
- o Configuration
- o Routing Exchange
- o Quality of Service (QoS) Exchange
- o Security Exchange
- o Filtering Exchange
- o Encapsulation/Tunneling Exchange
- o NAT and Application-Level Gateways
- o Measurement and Accounting
- o Diagnostics
- o CE Redundancy or CE Failover

4.1.1. Association, Capability Discovery, and Information Exchange

Association is the first step of the ForCES protocol exchange in which capability discovery and exchange happens between one or more CEs and the FEs. ForCES assumes that CEs and FEs already have

sufficient information to begin communication in a secure manner. The ForCES protocol is only applicable after CEs and FEs have discovered each other. ForCES makes no assumption about whether discovery was performed using a dynamic protocol or merely static configuration. Some discussion about how this can occur can be found in Section 6.4 of this document.

During the association phase, CEs and FEs exchange capability information with each other. For example, the FEs express the number of interface ports they provide, as well as the static and configurable attributes of each port.

In addition to initial configuration, the CEs and FEs also exchange dynamic configuration changes using ForCES. For example, FEs asynchronously inform the CEs of an increase/decrease in available resources or capabilities on the FE.

4.1.2. Topology Information Exchange

In this context, topology information relates to how the FEs are interconnected with each other with respect to packet forwarding. Topology discovery is outside the scope of the ForCES protocol. An implementation can choose its own method of topology discovery (for example, it can use a standard topology discovery protocol or apply a static topology configuration policy). Once the topology is established, the ForCES protocol may be used to transmit the resulting information to the CEs.

4.1.3. Configuration

ForCES is used to perform FE configuration. For example, CEs set configurable FE attributes such as IP addresses, etc. for their interfaces.

4.1.4. Routing Exchange

ForCES may be used to deliver packet forwarding information resulting from CE routing calculations. For example, CEs may send forwarding table updates to the FEs, so that they can make forwarding decisions. FEs may inform the CEs in the event of a forwarding table miss. ForCES may also be used to configure Equal Cost Multi-Path (ECMP) capability.

4.1.5. QoS Capabilities Exchange and Configuration

ForCES may be used to exchange QoS capabilities between CEs and FEs. For example, an FE may express QoS capabilities to the CE. Such capabilities might include metering, policing, shaping, and queuing functions. The CE may use ForCES to configure these capabilities.

4.1.6. Security Exchange

ForCES may be used to exchange security information between a CE and the FEs it controls. For example, the FE may use ForCES to express the types of encryption that it is capable of using in an IP Security (IPsec) tunnel. The CE may use ForCES to configure such a tunnel. The CEs would be responsible for the NE dynamic key exchanges and updates.

4.1.7. Filtering Exchange and Firewalls

ForCES may be used to exchange filtering information. For example, FEs may use ForCES to express the filtering functions, such as classification and action, that they can perform, and the CE may configure these capabilities.

4.1.8. Encapsulation/Tunneling Exchange

ForCES may be used to exchange encapsulation capabilities of an FE, such as tunneling, and the configuration of such capabilities.

4.1.9. NAT and Application-Level Gateways

ForCES may be used to exchange configuration information for Network Address Translators. Whilst ForCES is not specifically designed for the configuration of application-level gateway functionality, this may be in scope for some types of application-level gateways.

4.1.10. Measurement and Accounting

ForCES may be used to exchange configuration information regarding traffic measurement and accounting functionality. In this area, ForCES may overlap somewhat with functionality provided by network management mechanisms such as the Simple Network Management Protocol (SNMP). In some cases, ForCES may be used to convey information to the CE to be reported externally using SNMP. A further discussion of this capability is covered in Section 6 of this document.

4.1.11. Diagnostics

ForCES may be used for CEs and FEs to exchange diagnostic information. For example, an FE can send self-test results to a CE.

4.1.12. Redundancy and Failover

The ForCES architecture includes mechanisms that allow for multiple redundant CEs and FEs in a ForCES NE. The ForCES-model LFB definitions provide sufficient component details via component identifiers to be universally unique within an NE. The ForCES protocol includes mechanisms to facilitate transactions as well as atomicity across the NE.

Given the above, it is possible to deploy redundant CEs and FEs that incorporate failover.

4.2. CE-FE Link Capability

When using ForCES, the bandwidth of the CE-FE link is a consideration, and cannot be ignored. For example, sending a full routing table is reasonable over a high-bandwidth link, but could be non-trivial over a lower-bandwidth link. ForCES should be sufficiently future-proof to be applicable in scenarios where routing tables grow to several orders of magnitude greater than their current size. However, we also note that not all IP routers need full routing tables.

4.3. CE/FE Locality

ForCES is intended for environments where one of the following applies:

- o The control interconnect is some form of local bus, switch, or LAN, where reliability is high, closely controlled, and not susceptible to external disruption that does not also affect the CEs and/or FEs.
- o The control interconnect shares its fate with the FE's forwarding function. Typically this is because the control connection is also the FE's primary packet forwarding connection, and so if that link goes down, the FE cannot forward packets anyway.

The key guideline is that the reliability of the device should not be significantly reduced by the separation of control and forwarding functionality.

Taking this into account, ForCES is applicable in the following CE/FE localities:

Single Box NE:

chassis with multiple CEs and FEs set up. ForCES is applicable in localities consisting of control and forwarding elements that are components in the same physical box.

Example: a network element with a single control blade, and one or more forwarding blades, all present in the same chassis and sharing an interconnect such as Ethernet or Peripheral Component Interconnect (PCI). In this locality, the majority of the data traffic being forwarded typically does not traverse the same links as the ForCES control traffic.

Multiple Box NE:

separated CE and FE, where physical locality could be the same rack, room, or building; or long distances that could span across continents and oceans. ForCES is applicable in localities consisting of control and forwarding elements that are separated by a single hop or multiple hops in the network.

5. Security Considerations

The ForCES protocol allows for a variety of security levels [RFC5810]. When operating under a secured physical environment, or for other operational concerns (in some cases, performance issues), the operator may turn off all the security functions between CEs and FEs. When the operator makes a decision to secure the path between the FEs and CEs, then the operator chooses from one of the options provided by the TML. Security choices provided by the TML take effect during the pre-association phase of the ForCES protocol. An operator may choose to use all, some, or none of the security services provided by the TML in a CE-FE connection. A ForCES NE is required to provide CE/FE node authentication services, and may provide message integrity and confidentiality services. The NE may provide these services by employing IPsec or Transport Layer Security (TLS), depending on the choice of TML used in the deployment of the NE.

6. ForCES Manageability

From the architectural perspective, the ForCES NE is a single network element. As an example, if the ForCES NE is specifically a router that needs to be managed, then it should be managed in essentially the same way any router should be managed. From another perspective, element management could directly view the individual entities and interfaces that make up a ForCES NE. However, any element management

updates made directly on these entities and interfaces may compromise the control relationship between the CEs and the FEs, unless the update mechanism has been accounted for in the model used by the NE.

6.1. The NE as an Atomic Element

From the ForCES Requirements [RFC3654], Section 4, point 4:

A NE MUST support the appearance of a single functional device.

As a single functional device, a ForCES NE runs protocols, and each of the protocols has its own existing manageability aspects that are documented elsewhere. As an example, a router would also have a configuration interface. When viewed in this manner, the NE is controlled as a single routing entity, and no new management beyond what is already available for routers and routing protocols would be required for a ForCES NE. Management commands on a management interface to the NE will arrive at the CE and may require ForCES interactions between the CE and FEs to complete. This may impact the atomicity of such commands and may require careful implementation by the CE.

6.2. The NE as Composed of Manageable Elements

When viewed as a decomposed set of elements from the management perspective, the ForCES NE is divided into a set of one or more control elements, forwarding elements, and the interfaces between them. The interface functionality between the CE and the FE is provided by the ForCES protocol. A MIB module is provided for the purpose of gaining management information on the operation of the protocol described in Section 6.3 of this document.

Additionally, the architecture makes provisions for configuration control of the individual CEs and FEs. This is handled by elements called the FE Manager (FEM) and the CE Manager (CEM). Specifically, from the ForCES Requirements RFC [RFC3654], Section 4, point 4:

However, external entities (e.g., FE Managers and CE Managers) MAY have direct access to individual ForCES protocol elements for providing information to transition them from the pre-association to the post-association phase.

6.3. ForCES Protocol MIB

The ForCES MIB [RFC5813] defines a primarily read-only MIB module that captures information related to the ForCES protocol. This includes state information about the associations between CE(s) and FE(s) in the NE.

The ForCES MIB does not include information that is specified in other MIB modules, such as packet counters for interfaces, etc.

More specifically, the information in the ForCES MIB module relative to associations includes:

- o identifiers of the elements in the association
- o state of the association
- o configuration parameters of the association
- o statistics of the association

6.3.1. MIB Management of an FE

While it is possible to manage an FE from an element manager, several requirements relating to this have been included in the ForCES Requirements.

From the ForCES Requirements [RFC3654], Section 4, point 14:

1. The ability for a management tool (e.g., SNMP) to be used to read (but not change) the state of FE SHOULD NOT be precluded.
2. It MUST NOT be possible for management tools (e.g., SNMP, etc) to change the state of a FE in a manner that affects overall NE behavior without the CE being notified.

The ForCES Framework [RFC3746], Section 5.7, goes further in discussing the manner in which FEs should handle management requests that are specifically directed to the FE:

(For a ForCES NE that is an IP router,) RFC 1812 [RFC1812] also dictates that "Routers must be manageable by SNMP". In general, for the post-association phase, most external management tasks (including SNMP) should be done through interaction with the CE in order to support the appearance of a single functional device. Therefore, it is recommended that an SNMP agent be implemented by CEs and that the SNMP messages received by FEs be redirected to their CEs. AgentX framework defined in RFC 2741 [RFC2741]) may be applied here such that CEs act in the role of master agent to process SNMP messages while FEs act in the role of subagent to provide access to the MIB objects residing on FEs. AgentX protocol messages between the master agent (CE) and the subagent (FE) are encapsulated and transported via ForCES, just like data packets from any other application layer protocols.

6.4. The FEM and CEM

Though out of scope for the initial ForCES specification effort, the ForCES architecture includes two entities: the CE Manager (CEM) and the FE Manager (FEM). From the ForCES Protocol Specification [RFC5810]:

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.

FE Manager (FEM):

A logical entity responsible for generic FE management tasks. It is used during the pre-association phase to determine with which CE(s) an FE should communicate.

7. Contributors

Mark Handley was an initial author involved in the earlier versions of this document.

8. Acknowledgments

Many of the participants in the ForCES WG, as well as fellow employees of the authors, have provided valuable input into this work. Particular thanks go to Jamal Hadi Salim, our WG chair and document shepherd; and to Adrian Farrel, the AD for the area; for their review, comments, and encouragement, without which this document might never have been completed.

9. References

9.1. Normative References

- [RFC1812] Baker, F., "Requirements for IP Version 4 Routers", RFC 1812, June 1995.
- [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.

9.2. Informative References

- [RFC2741] Daniele, M., Wijnen, B., Ellison, M., and D. Francisco, "Agent Extensibility (AgentX) Protocol Version 1", RFC 2741, January 2000.
- [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.

Authors' Addresses

Alan Crouch
Intel
2111 NE 25th Avenue
Hillsboro, OR 97124
USA

Phone: +1 503 264 2196
EMail: alan.crouch@intel.com

Hormuzd Khosravi
Intel
2111 NE 25th Avenue
Hillsboro, OR 97124
USA

Phone: 1-503-264-0334
EMail: hormuzd.m.khosravi@intel.com

Avri Doria (editor)
LTU
Lulea University of Technology
Sweden

Phone: +46 73 277 1788
EMail: avri@acm.org

Xin-ping Wang
Huawei
Beijing
China

Phone: +86 10 82836067
EMail: carly.wang@huawei.com

Kentaro Ogawa
NTT Corporation
3-9-11 Midori-cho
Musashino-shi, Tokyo 180-8585
Japan

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