forces
Internet-Draft
Intended status: Informational
Expires: April 13, 2010

A. Crouch
H. Khosravi
Intel
A. Doria
LTU
X. Wang
Huawei
K. Ogawa
NTT Corporation

October 10, 2009

ForCES Applicability Statement
draft-ietf-forces-applicability-07

Status of this Memo

Copyright Notice

Abstract

The 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.

Table of Contents

1. Purpose

The purpose of the ForCES Applicability Statement is to capture the intent of the ForCES protocol [I-D.ietf-forces-protocol] designers as to how the protocol could be used (in conjunction with the ForCES model [I-D.ietf-forces-model]).

2. Overview

The 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.

3. Terminology

A set of terminology associated with ForCES is defined in [3, 4]. That terminology is reused here and the reader is directed to [3, 4] for the following definitions:

- CE: Control Element.

- FE: Forwarding Element.

- ForCES: ForCES protocol.

- TML: Transport Mapping Layer.

4. Applicability to IP Networks

The purpose of this section is to list the areas of ForCES applicability in IP network devices. Relatively low end routing systems may be implemented on simple hardware which performs both control and packet forwarding functionality. ForCES may not make sense for such devices.

Higher end routing systems typically distribute work amongst interface processing elements, and these devices (FEs) therefore need to communicate with the control element(s) to perform their job. ForCES provides a standard way to do this communication.

The remainder of this section lists the applicable services which 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 Discovery, Capability Information Exchange

o Topology Information Exchange

o Configuration

o Routing Exchange

o 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. Discovery, Capability Information Exchange

Discovery is the process by which CEs and FEs learn of each other's existence. ForCES assumes that CEs and FEs already know sufficient information to begin communication in a secure manner. The ForCES protocol is only applicable after CEs and FEs have found each other. ForCES makes no assumption about whether discovery was performed using a dynamic protocol or merely static configuration.

During the discovery 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 CE 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 use a standard topology discovery protocol like LLDP, BFD;or apply a static topology configuration policy).Once the topology is established, ForCES protocol may be used to transmit the resulting information to the CE.

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 CE in the event of a forwarding table miss.

4.1.5. QoS Exchange

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 CEs and FEs. For example, the FE may use ForCES to express the types of encryption that it is capable of using in an IPsec tunnel. The CE may use ForCES to configure such a tunnel.

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 alternative network management mechanisms such as SNMP. In some cases ForCES may be used to convey information to the CE to be reported externally using SNMP.

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 the CE.

4.1.12. CE Redundancy or CE Failover

CE failover and redundancy are out of scope in the initial version of ForCES protocol. Basic mechanisms for CE redundancy/failover are not presently implemented. Broad concepts such as implementing CE Redundancy, CE Failover, and CE-CE communication, while not precluded by the ForCES architecture, are considered outside the scope of ForCES protocol. ForCES protocol is designed to handle CE- FE communication, and is not intended for CE-CE communication.

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 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:

o single box NE: chassis with multiple CEs and FEs setup. ForCES is applicable in localities consisting of control and forwarding elements which 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 PCI. In this locality, the majority of the data traffic being forwarded typically does not traverse the same links as the ForCES control traffic.

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

5. Security Considerations

The ForCES architecture allows for a variety of security levels[6].

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 CE and FE. When the operator makes a decision to secure the path between the FE and CE 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 confidentially services. The NE may provide these services by employing IPSEC or TLS depending on the choice of TML used in the deployment of the NE.

6. ForCES Manageability

From the management perspective, an NE can be viewed in at least two ways. From one perspective, it is a single network element, specifically a router that needs to be managed in essentially the same way any router is managed. From another perspective element management can view the individual entities and interfaces that make up a ForCES NE.

6.1. NE as an atomic element

From the ForCES requirements RFC 3654, 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 it own existing manageability aspects that are documented elsewhere. As a router it 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.

6.2. 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 of 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. As with all IETF protocols a MIB is provided for the purposes of managing the protocol.

Additionally the architecture makes provision for configuration

control of the individual CEs and FEs. This is handled by elements named FE manager (FEM) and the CE manager (CEM). Specifically from the ForCES requirements RFC [RFC 3654], 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 post-association phase.

6.3. ForCES Protocol MIB

The ForCES MIB [I-D.ietf-forces-mib] is a primarily read-only MIB 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 MIBs, such as packet counters for interfaces, etc.

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

- identifiers of the elements in the association

- state of the association

- configuration parameters of the association

- statistics of the association

6.3.1. MIB Management of an FE

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

From the ForCES Requirements [RFC 3654], 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 Requirements [RFC 3654], Section 5.7, goes further in discussing the manner in which FEs should handle management requests that are specifically directed to the FE:

RFC 1812 [2] 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 ([6]) may be applied here such that CEs act in the role of master agent to process SNMP protocol 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 include two entities, the CE Manager (CEM) and the FE Manager (FEM). From the ForCES Protocols Specification [I-D.ietf-forces-protocol].

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 pre-association phase to determine with which CE(s) an FE should communicate.

7. Contributors

The following are the contributors who were instrumental in the creation of earlier releases of this document or who gave good suggestions to this document.

Mark Handley,ICIR.

8. IANA Considerations

This document has no IANA actions.

[RFC Editor: please remove this section prior to publication.]

9. Acknowledgments

Many of the colleagues in our companies and participants in the ForCES mailing list have provided invaluable input into this work. Particular thanks to Jamal Hadi Salim.

10. References

10.1. Normative References

[I-D.ietf-forces-mib]
HAAS, R., "ForCES MIB", draft-ietf-forces-mib-10 (work in progress), September 2008.

[I-D.ietf-forces-model]
Halpern, J. and J. Salim, "ForCES Forwarding Element Model", draft-ietf-forces-model-16 (work in progress), October 2008.

[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.

[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.

10.2. Informative References

[RFC3015] Cuervo, F., Greene, N., Rayhan, A., Huitema, C., Rosen, B., and J. Segers, "Megaco Protocol Version 1.0", RFC 3015, November 2000.

[RFC3292] Doria, A., Hellstrand, F., Sundell, K., and T. Worster, "General Switch Management Protocol (GSMP) V3", RFC 3292, June 2002.

Authors' Addresses

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

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

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

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

Avri Doria
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