Internet Engineering Task Force
INTERNET-DRAFT
draft-ietf-forces-applicability-00.txt

ForCES WG
Alan Crouch/Intel Labs
Mark Handley/ACIRI

19 June 2002
Expires: December 2002

# ForCES Applicability Statement

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of Section 10 of RFC2026.

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.

## Abstract

The ForCES protocol defines a standard framework and mechanism for the interconnection between Control Elements and Forwarding Engines 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.

## 1. Purpose

The purpose of the ForCES Applicability Statement is to capture the intent of the ForCES protocol designers as to how the protocol should be used. This Applicability Statement will evolve alongside the protocol, with the intent that it be published as an informational RFC around the same time that the ForCES protocol is published as a standards-track RFC.

## 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 [1]. That terminology is reused here and the reader is directed to [1] for the following definitions:

- CE: Control Element.
- FE: Forwarding Element.
- ForCES: ForCES protocol.

## 4. Applicability to IP Networks

The purpose of this section is to list the areas of ForCES applicability in IP network devices. Relatively low performance devices may be implemented on a simple processor which performs both control and packet forwarding functionality. ForCES is not applicable for such devices. Higher performance devices typically distribute work amongst interface processors, 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:

- Discovery, Capability Information Exchange
- Topology Information Exchange
- Configuration
- Routing Exchange
- QoS Exchange
- Security Exchange
- Filtering Exchange
- Encapsulation/Tunneling Exchange
- NAT and Application-level Gateways
- Measurement and Accounting
- Diagnostics
- 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 may exchange capability information with each other. For example, the FEs may 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 may also exchange dynamic configuration changes using ForCES. For example, FE's 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. Whilst topology discovery is outside the scope of the ForCES protocol, a standard topology discovery protocol may be selected and used to "learn" the topology, and then the 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.

#### 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. However, in other cases it may make more sense for the FE to directly speak SNMP.

#### 4.1.11. Diagnostics

ForCES may be used for CE's and FE's to exchange diagnostic information. For example, an FE can send self-test results to the CE.

#### 4.1.12. CE Redundancy or CE Failover

ForCES is a master-slave protocol where FE's are slaves and CE's are masters. Basic mechanisms for CE redundancy/failover are provided in ForCES protocol. 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 Capacity

When using ForCES, the bandwidth of the CE-FE link is a consideration, and cannot be ignored. For example, sending a full routing table of 110K routes is reasonable over a 100Mbit Ethernet interconnect, 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 (approximately 100K routes). However, we also note that not all IP routers need full routing tables.

### 4.3. CE/FE Locality

We do not intend ForCES to be applicable in configurations where the CE and FE are located arbitrarily in the network. In particular, ForCES is intended for environments where one of the following applies:

- 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.

- 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.

ForCES is applicable in localities consisting of control and forwarding elements which are either components in the same physical box, or are separated at most by one local network hop (historically referred to as "Very Close" localities).

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.

## 5. Limitations and Out-of-Scope Items

ForCES was designed to enable logical separation of control and forwarding planes in IP network devices. However, ForCES is not

intended to be applicable to all services or to all possible CE/FE localities.

The purpose of this section is to list limitations and out-of-scope items for ForCES.

## 5.1. Out of Scope Services

The following control-forwarding plane services are explicitly not addressed by ForCES:

- Label Switching
- Multimedia Gateway Control (MEGACO).

### 5.1.1. Label Switching

Label Switching is the purview of the GSMP Working Group in the Sub- IP Area of the IETF. GSMP is a general purpose protocol to control a label switch. GSMP defines mechanisms to separate the label switch data plane from the control plane label protocols such as LDP [5]. For more information on GSMP, see [4].

### 5.1.2. Separation of Control and Forwarding in Multimedia Gateways"

MEGACO defines a protocol used between elements of a physically decomposed multimedia gateway. Separation of call control channels from bearer channels is the purview of MEGACO. For more information on MEGACO, see [7].

## 5.2. Localities

The ForCES protocol was intended to work within the localities described in section 4.3. While the ForCES protocol might be able to work in a wider range of circumstances, anyone trying to do so should be aware that it has not been designed or evaluated for such use. In particular, there are many clear cases where arbitrary separation of control and forwarding would render network operation significantly more fragile than non-separated forwarding.

Examples of localities where ForCES was not designed or evaluated for use are:

- Localities where there are multiple hops between CE and FE.

- Localities where hops between the CE and FE are dynamically routing using IP routing protocols.

- Localities where the loss of the CE-FE link is of non-negligible probability.

- Localities where two or more FEs controlled by the same CE cannot communicate, either directly, or indirectly via other FEs controlled by the same CE.

## 6. Security Considerations

The security of ForCES protocol will be addressed in the Protocol Specification [2]. For security requirements, see architecture requirement #5 and protocol requirement #2 in the Requirements Draft [1]. The ForCES protocol assumes that the CE and FE are in the same administration, and have shared secrets as a means of administration. Whilst it might be technically feasible to have the CE and FE administered independently, we strongly discourage such uses, because they would require a significantly different trust model from that ForCES assumes.

## 7. Normative

[1] Anderson, T et. al., "Requirements for Separation of IP Control and Forwarding", draft-ietf-forces-requirements-01.txt, Intel Labs, September 2001.

[2] ForCES Protocol Specification (to-be-written)

## 8. Informative

[3] Salim, J e. al., "Netlink as an IP Services Protocol", draft-salim-netlink-jhsk-01.txt, Znyx Networks, September 2001.

[4] Doria, A, Sundell, K, Hellstrand, F, Worster, T, "General switch Management Protocol V3," Internet Draft draft-ietf-gsmp-06.txt, July **2000. work in progress**

[5] Andersson et al., "LDP Specification" RFC 3036, January 2001

[6] Bradner, S, "Key words for use in RFCs to Indicate Requirement Levels", RFC 2119, Harvard University, March 1997.

[7] F. Cuervo et al., "Megaco Protocol Version 1.0" RFC 3015, November

2000

## 9. Acknowledgments

The authors wish to thank Jamal Hadi, Hormuzd Khosravi, Vip Sharma, and many others for their invaluable contributions.

## 10. Author's Addresses

Alan Crouch
Intel Labs
2111 NE 25th Avenue
Hillsboro, OR 97124 USA
Phone: +1 503 264 2196
Email: alan.crouch@intel.com

Mark Handley
ICSI
1947 Center Street, Suite 600
Berkeley, CA 94704, USA
Email: mjh@icsi.berkeley.edu