Hormuzd Khosravi
Intel Corp.
Furquan Ansari
Lucent Tech.
Jon Maloy
Ericsson

Internet Draft
Document: draft-khosravi-forces-tcptml-00.txt
Expires: May 2005
Working Group: ForCES

November 2004

TCP/IP based TML (Transport Mapping Layer) for ForCES protocol

Status of this Memo

By submitting this Internet-Draft, I certify that any applicable patent or other IPR claims of which I am aware have been disclosed, and any of which I become aware will be disclosed, in accordance with RFC 3668.

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.

Conventions used in this document

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 [2].

Abstract

This document defines the TCP/IP based TML (Transport Mapping Layer) for the ForCES protocol. It explains the rationale for choosing the transport protocols and also describes how this TML addresses all

the requirements described in the Forces [3] requirements and ForCES protocol [7] document.

Table of Content

1. Definitions

The following definitions are taken from [3], [5]

ForCES Protocol - While there may be multiple protocols used within the overall ForCES architecture, the term "ForCES protocol" refers only to the protocol used at the Fp reference point in the ForCES Framework in RFC3746 [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.

ForCES Protocol Layer (ForCES PL) -- A layer in ForCES protocol architecture that defines the ForCES protocol messages, the protocol state transfer scheme, as well as the ForCES protocol architecture

itself (including requirements of ForCES TML (see below)). Specifications of ForCES PL are defined by this document.

ForCES Protocol Transport Mapping Layer (ForCES TML) -- A layer in ForCES protocol architecture that specifically addresses the 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. This document defines a TCP/IP based ForCES TML.

2. Introduction

The ForCES (Forwarding and Control Element Separation) working group in the IETF is defining the architecture and protocol for separation of control and forwarding elements in network elements such as routers. [3], [4] define both architectural and protocol requirements for the communication between CE and FE. The ForCES protocol layer [7] describes the protocol specification. It is envisioned that the ForCES protocol would be independent of the interconnect technology between the CE and FE and can run over multiple transport technologies and protocol. Thus a Transport Mapping Layer has been defined in the protocol framework that will take care of mapping the protocol messages to specific transports. This document defines the TCP/IP based TML for the ForCES protocol layer. It also addresses all the requirements for the TML including security, reliability, etc.

3. Protocol Framework Overview

The reader is referred to the Framework document [RFC3746], and in particular sections 3 and 4, for architectural overview and where and how the ForCES protocol fits in. There may be some content overlap between the ForCES protocol draft [7] and this section in order to provide clarity.

The ForCES protocol constitutes two pieces: the PL and TML layer. This is depicted in Figure 1 below.

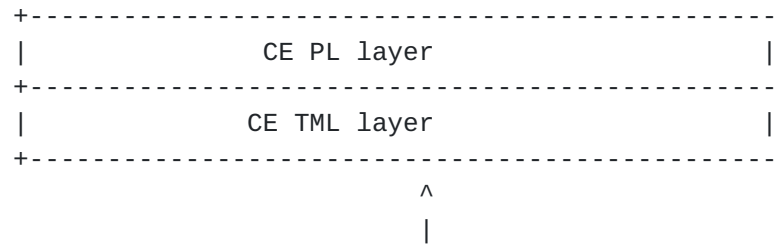

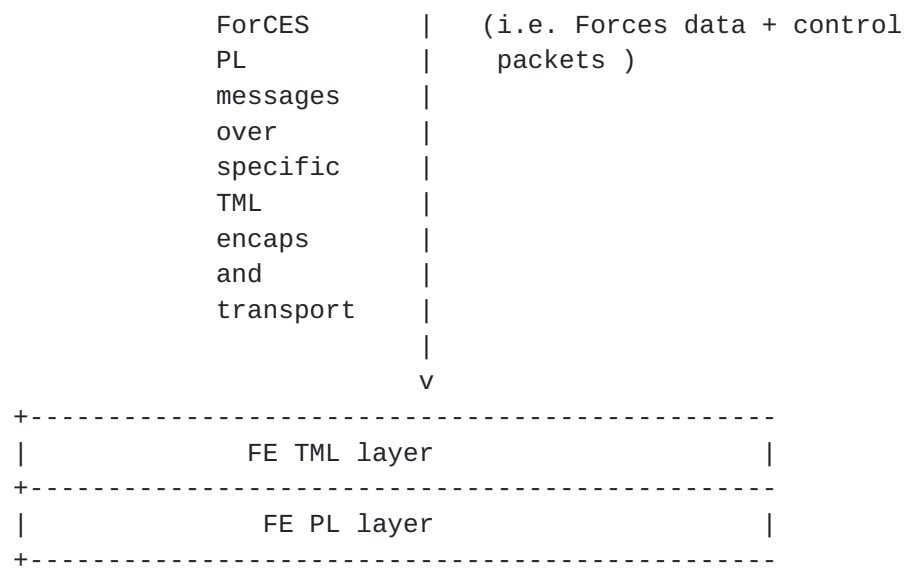

Figure 1: ForCES Interface

The PL layer is in fact the ForCES protocol. Its semantics and message layout are defined in [7]. The TML Layer is necessary to connect two ForCES PL layers as shown in Figure 1 above.

Both the PL and TML layers are standardized by the IETF. While only one PL layer is defined, different TMLs are expected to be standardized. To interoperate the TML layer at the CE and FE are expected to be of the same definition.

On transmit, the PL layer delivers its messages to the TML layer. The TML layer delivers the message to the destination TML layer(s). On reception, the TML delivers the message to its destination PL layer(s).

3.1.1 The PL layer

The PL is common to all implementations of ForCES and is standardized by the IETF [7]. The PL layer is responsible for associating an FE or CE to an NE. It is also responsible for tearing down such associations. An FE uses the PL layer to throw various subscribed-to events to the CE PL layer as well as respond to various status requests issued from the CE PL. The CE configures both the FE and associated LFBs attributes using the PL layer. In addition the CE may send various requests to the FE to activate or deactivate it, reconfigure its HA parameterization, subscribe to specific events etc.

3.1.2 The TML layer

The TML layer is essentially responsible for transport of the PL layer messages. The TML is where the issues of how to achieve transport level reliability, congestion control, multicast, ordering, etc. are handled. It is expected more than one TML will be standardized. The different TMLs each could implement things differently based on 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 should be portable across all TMLs, because all TMLs have the same top edge semantics.

4. TCP/IP TML Overview

The TCP/IP TML consists of two TCP connections between the CE and FE over which the protocol messages are exchanged. One of the connection is called the Control channel, over which control messages are exchanged, the other is called data channel over which external protocol packets, such as routing packets will be exchanged. The TCP connections will use unique server port numbers for each of the channels. In addition to this, this TML will provide mechanisms to prioritize the messages over the different channels.

Some of the rationale of choosing this transport mechanism as well as explanation of how it meets the TML requirements is explained below.

4.1.Rationale for using TCP/IP

TCP meets all the reliability requirements (no losses, no data corruption, no re-ordering of data) for the ForCES protocol/TML and also provides congestion control mechanism, which is important to meet the scalability requirement. In addition, it helps with interoperability since TCP is a well-understood, widely deployed transport protocol. Using TCP also enables this TML and the protocol to work seamlessly in single hop and multihop scenarios.

4.2.Separate Control and Data channels

The ForCES NEs are subject to Denial of Service (DoS) attacks [Requirements Section 7 #15]. A malicious system in the network can flood a ForCES NE with bogus control packets such as spurious RIP or OSPF packets in an attempt to disrupt the operation of and the communication between the CEs and FEs. In order to protect against this situation, the TML uses separate control and data channels for communication between the CEs and FEs.

The data channel carries the control protocol packets such as RIP, OSPF messages as outlined in Requirements [3] Section 7 #10, which are carried in ForCES Packet Redirect messages [7], between the CEs and FEs. All the other ForCES messages, which are used for configuration/capability exchanges, event notification, etc, are carried over the control channel. The data channel is set up only after the control channel is set up and the capability exchange has successfully taken place between the FEs and CEs. The CE signals the FE to establish the data channel at the appropriate time and provides it with the channel addressing information, such as, port number in case of TCP (see section 5.3). By default, the data channel is established on the CE control channel port number +1.

The reliability requirements for the data channel messages are different from that of the control messages [Reqs] i.e. they dont require strict reliability in terms of retransmission, etc. However congestion control is important for the data channel because in case of DoS attacks, if an unreliable transport such as UDP is used for the data traffic, it can more easily overflow the physical connection, overwhelming the control traffic with congestion. Thus we need a transport protocol that provides congestion control but does not necessarily provide full reliability. Datagram Congestion Control Protocol (DCCP) [11], which is currently being defined, is a transport protocol that exactly meets this requirement. However since it is currently not an IETF standard RFC, and the authors are unaware of any existing implementations, this TML uses TCP as transport protocol for the data channel (for IP interconnect). TCP provides the congestion control mechanism required for the data channel and its wide deployment eases interoperability.

4.3.Reliability

TCP provides the reliability (no losses, no data corruption, no re-ordering of data) required for ForCES protocol control messages.

4.4.Congestion Control

TCP provides congestion control needed to satisfy this requirement.

4.5.Security

This TML uses TLS [8] to provide security in insecure environments. Please see section 6 on security considerations for more details.

4.6.Addressing

This TML uses addressing provided by IP layer.

4.7.Prioritization

This TML provides prioritization of messages sent over control channel as compared to the data channel. This has also been found to be useful in face of DoS attacks on the protocol. The details of this are TBD.

4.8.HA Decisions

TBD

4.9.Encapsulations Used

Other than the TCP/IP header, no other encapsulations will be added to the ForCES protocol messages.

5. Example Scenarios

5.1.Establishment of Association

TBD

5.2.Steady State Communication

TBD

6. Security Considerations

If the CE or FE are in a single box and network operator is running under a secured environment then it is up to the network administrator to turn off all the security functions. This is configured during the pre-association phase of the protocol.

When the CEs, FEs are running over IP networks or in an insecure environment, this TML uses TLS [8] to provide security. The security association between the CEs and FEs MUST be established before any ForCES protocol messages are exchanged between the CEs and FEs.

6.1.TLS Usage for this TML

This section is applicable for CE or FE endpoints that use the TML with TLS [8] to secure communication.

Since CE is master and FEs are slaves, the FEs are TLS clients and CEs are TLS server. The endpoints that implement TLS MUST perform mutual authentication during TLS session establishment process. CE must request certificate from FE and FE needs to pass the requested information.

We recommend TLS-RSA-with-AES-128-CBC-SHA¶ cipher suite, but CE or FE may negotiate other TLS cipher suites. TLS must be used for all control channel messages. TLS is optional for the data channel since data channel packets are not encrypted externally to the NE.

This TML uses TLS to provide security when the NE is in an insecure environment. This is because IPsec provides less flexibility when configuring trust anchors since it is transparent to the application and use of Port identifiers is prohibited within IKE Phase 1. This provides restriction for IPsec to configure trust anchors for each application separately and policy configuration is common for all applications.

7. IANA Considerations

The TCP/IP TML needs to have a two well-defined TCP port numbers, which needs to be assigned by IANA.

8. References

8.1.Normative References

1. S. Bradner, "The Internet Standards Process -Revision 3", RFC 2026, October 1996.

2. S. Bradner, "Keywords for use in RFCs to Indicate Requirement Levels", RFC2119 (BCP), IETF, March 1997.

3. Khosravi, et al., ¶Requirements for Separation of IP Control and Forwarding¶, RFC 3654, November 2003.

4. L. Yang, et al., ¶ ForCES Architectural Framework¶, RFC 3746, April 2004.

5. L. Yang, et al., ¶ ForCES Forwarding Element Functional Model¶, work in progress¶, July 2004,<draft-ietf-forces-model-03.txt>

6. A. Audu, et al., Forwarding and Control Element protocol (FACT)" draft-gopal-forces-fact-06.txt, February 2004.

7. A. Doria, et al., ¶ForCES protocol specification¶, draft-ietf-forces-protocol-00.txt, September 2004.

8.2.Informative References

8. Dierks, T., Allen, C., Treese, W., Karlton, P., Freier, A. and P. Kocher, "The TLS Protocol Version 1.0", RFC 2246, January 1999.

9. Jungmaier, A., Rescorla, E. and M. Tuexen, "Transport Layer Security over Stream Control Transmission Protocol", RFC 3436, December 2002.

10. R. Stewart, et al., Stream Control Transmission Protocol (SCTP)¶, RFC 2960, October 2000.

11. E. Kohler, M. Handley, S. Floyd, J. Padhye, Datagram Congestion Control Protocol (DCCP)¶, draft-ietf-dccp-spec-04.txt, June 2003.

12. Floyd, S., Congestion Control Principles¶, RFC 2914, September 2000.

13. A. Doria, F. Hellstrand, K. Sundell, T. Worster, General Switch Management Protocol (GSMP) V3¶, RFC 3292, June 2002.

14. H. Balakrishnan, et al. The Congestion Manager¶, RFC 3124, June 2001.

9. Acknowledgments

10. Authors' Addresses

Hormuzd Khosravi
Intel
2111 NE 25th Avenue
Hillsboro, OR 97124
Phone: 1-503-264-0334
Email: hormuzd.m.khosravi@intel.com

Furquan Ansari
101 Crawfords Corner Road
Holmdel, NJ 07733
USA

Phone: +1 732-949-5249
Email: furquan@lucent.com

Jon Maloy
Ericsson Research Canada
8400 Boul Decarie
Ville Mont-Royal, Quebec H4P 2N2
Canada
Phone: 1-514-345-7900
Email: jon.maloy@ericsson.com