

Internet Engineering Task Force  
Internet-Draft  
Intended status: Informational  
Expires: September 18, 2014

J. Hadi Salim  
Mojatatu Networks  
March 17, 2014

ForCES Protocol Extensions  
draft-ietf-forces-protoextension-01

## Abstract

Experience in implementing and deploying ForCES architecture has demonstrated need for a few small extensions both to ease programmability and to improve wire efficiency of some transactions. This document describes extensions to the ForCES Protocol Specification[RFC 5810] semantics to achieve that end goal.

## Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on September 18, 2014.

## Copyright Notice

Copyright (c) 2014 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.

Internet-Draft

ForCES Protocol Extensions

March 2014

Table of Contents

- [1. Terminology and Conventions . . . . .](#) [3](#)
- [1.1. Requirements Language . . . . .](#) [3](#)
- [1.2. Definitions . . . . .](#) [3](#)
- [2. Introduction . . . . .](#) [4](#)
- [3. Problem Overview . . . . .](#) [5](#)
- [3.1. Table Ranges . . . . .](#) [5](#)
- [3.2. Error codes . . . . .](#) [5](#)
- [4. Protocol Update Proposal . . . . .](#) [6](#)
- [4.1. Table Ranges . . . . .](#) [6](#)
- [4.2. Error Codes . . . . .](#) [7](#)
- [4.2.1. New Codes . . . . .](#) [8](#)
- [4.2.2. Vendor Codes . . . . .](#) [8](#)
- [4.2.3. Extended Result TLV . . . . .](#) [8](#)
- [4.3. Large Object Dumping . . . . .](#) [9](#)
- [5. Acknowledgements . . . . .](#) [10](#)
- [6. IANA Considerations . . . . .](#) [10](#)
- [7. Security Considerations . . . . .](#) [10](#)
- [8. References . . . . .](#) [11](#)
- [8.1. Normative References . . . . .](#) [11](#)
- [8.2. Informative References . . . . .](#) [11](#)
- [Author's Address . . . . .](#) [11](#)

## 1. Terminology and Conventions

### 1.1. Requirements Language

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](#)].

### 1.2. Definitions

This document reiterates the terminology defined by the ForCES architecture in various documents for the sake of clarity.

**FE Model** - The FE model is designed to model the logical processing functions of an FE. The FE model proposed in this document includes three components; the LFB modeling of individual Logical Functional Block (LFB model), the logical interconnection between LFBs (LFB topology), and the FE-level attributes, including FE capabilities. The FE model provides the basis to define the information elements exchanged between the CE and the FE in the ForCES protocol [[RFC5810](#)].

**LFB (Logical Functional Block) Class (or type)** - A template that represents a fine-grained, logically separable aspect of FE processing. Most LFBs relate to packet processing in the data path. LFB classes are the basic building blocks of the FE model.

**LFB Instance** - As a packet flows through an FE along a data path, it flows through one or multiple LFB instances, where each LFB is an instance of a specific LFB class. Multiple instances of the same LFB class can be present in an FE's data path. Note that we often refer to LFBs without distinguishing between an LFB class and LFB instance when we believe the implied reference is obvious for the given context.

**LFB Model** - The LFB model describes the content and structures in

an LFB, plus the associated data definition. XML is used to provide a formal definition of the necessary structures for the modeling. Four types of information are defined in the LFB model. The core part of the LFB model is the LFB class definitions; the other three types of information define constructs associated with and used by the class definition. These are reusable data types, supported frame (packet) formats, and metadata.

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, but not how the per-packet state is

implemented within actual hardware. Metadata is sent between the FE and the CE on redirect packets.

ForCES Component - A ForCES Component is a well-defined, uniquely identifiable and addressable ForCES model building block. A component has a 32-bit ID, name, type, and an optional synopsis description. These are often referred to simply as components.

LFB Component - An LFB component is a ForCES component that defines the Operational parameters of the LFBs that must be visible to the CEs.

ForCES Protocol - Protocol that runs in the Fp reference points in the ForCES Framework [[RFC3746](#)].

ForCES Protocol Layer (ForCES PL) - A layer in the ForCES protocol architecture that defines the ForCES protocol messages, the protocol state transfer scheme, and the ForCES protocol architecture itself as defined in the ForCES Protocol Specification [[RFC5810](#)].

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, ordering, etc. the ForCES SCTP TML [[RFC5811](#)] describes a TML that is mandated for ForCES.

## [2.](#) Introduction

Experience in implementing and deploying ForCES architecture has demonstrated need for a few small extensions both to ease programmability and to improve wire efficiency of some transactions. This document describes a few extensions to the ForCES Protocol Specification [[RFC5810](#)] semantics to achieve that end goal.

This document describes and justifies the need for 2 small extensions which are backward compatible. The document also clarifies on top of [[RFC5810](#)] how dumping of large components is achieved.

1. A table range operation to allow a controller or control application to request an arbitrary range of table rows.
2. Improved Error codes returned to the controller (or control application) to improve granularity of existing defined error

Hadi Salim

Expires September 18, 2014

[Page 4]

---

Internet-Draft

ForCES Protocol Extensions

March 2014

codes.

3. FE response to a GET request (example to a large table) may not fit in one PL message, for example due to limited TLV space.

## [3.](#) Problem Overview

In this section we present sample use cases to illustrate the challenge being addressed.

### [3.1.](#) Table Ranges

Consider, for the sake of illustration, an FE table with 1 million reasonably sized table rows which are sparsely populated. Assume, again for the sake of illustration, that there are 2000 table rows sparsely populated between the row indices 23-10023.

ForCES GET and DEL requests sent from a controller (or control app) are prepended with a path to a component and sent to the FE. In the case of indexed tables, the component path can either be to a table or a table row index. The approaches for retrieving or deleting a

sizeable number of table rows is at the programmatically level (from an application point of view) unfriendly, tedious, and abusive of both compute and bandwidth resources.

As an example, a control application attempting to retrieve the first 2000 table rows appearing between row indices 23 and 10023 can achieve its goal in one of:

- o Dump the whole table and filter for the needed 2000 table rows.
- o Send upto 10000 ForCES PL requests with monotonically incrementing indices and stop when the needed 2000 entries are retrieved.
- o If the application had knowledge of which table rows existed (not unreasonable given the controller is supposed to be aware of state within an NE), then the application could take advantage of ForCES batching to send fewer large messages (each with different path entries for a total of two thousand).

As argued, while the above options exist - all are tedious.

### [3.2.](#) Error codes

[RFC5810] has defined a generic set of error codes that are to be returned to the CE from an FE. Deployment experience has shown that it would be useful to have more fine grained error codes. As an

Hadi Salim

Expires September 18, 2014

[Page 5]

---

Internet-Draft

ForCES Protocol Extensions

March 2014

example, the error code E\_NOT\_SUPPORTED could be mapped to many FE error source possibilities that need to be then interpreted by the caller based on some understanding of the nature of the sent request. This makes debugging more time consuming.

## [4.](#) Protocol Update Proposal

This section describes proposals to update the protocol for issues discussed in [Section 3](#)

### [4.1.](#) Table Ranges

We propose to add a Table-range TLV (type ID 0x117) that will be associated with the PATH-DATA TLV in the same manner the KEYINFO-TLV

is.




Figure 1: ForCES table range request Layout

Figure 1 shows how this new TLV is constructed.

```
OPER = GET
PATH-DATA:
  flags = F_SELTABRANGE, IDCount = 2, IDs = {1,6}
  TABLERANGE-TLV content = {11,23}
```

Figure 2: ForCES table range request

Figure 2 illustrates a GET request for a range of rows 11 to 23 of a table with component path of "1/6".

Path flag of F\_SELTABRANGE (0x2 i.e bit 1, where bit 0 is F\_SELKEY as defined in [RFC 5810](#)) is set to indicate the presence of the TABLERANGE-TLV. The pathflag bit F\_SELTABRANGE can only be used in a GET or DEL and is mutually exclusive with F\_SELKEY. The FE MUST enforce those constraints and reject a request with an error code of E\_INVALID\_TFLAGS with a description of what the problem is (refer to [Section 4.2](#)).

The TABLERANGE-TLV contents constitute:

- o A 32 bit start index. An index of 0 implies the beginning of the table row.
- o A 32 bit end index. A value of 0xFFFFFFFFFFFFFFFF implies the last entry.

The response for a table range query will either be:

- o The requested table data returned (when at least one referenced row is available); in such a case, a response with a path pointing to the table and whose data content contain the row(s) will be sent to the CE. The data content MUST be encapsulated in sparsedata TLV. The sparse data TLV content will have the "I" (in ILV) for each table row indicating the table indices.
- o An EXTENDEDRESULT-TLV (refer to [Section 4.2.3](#)) when:
  - \* Response is to a range delete request. The Result will either be:
    - + A success if any of the requested-for rows is deleted
    - + A proper error code if none of the requested for rows can be deleted
  - \* data is absent where the result code of E\_EMPTY with an optional content string describing the nature of the error (refer to [Section 4.2](#)).
  - \* When both a path key and path table range are reflected on the the pathflags, an error code of E\_INVALID\_TFLAGS with an optional content string describing the nature of the error (refer to [Section 4.2](#)).
  - \* other standard ForCES errors (such as ACL constraints trying to retrieve contents of an unreadable table), accessing unknown components etc.

#### [4.2](#). Error Codes

We propose several things:

1. A new set of error codes.
2. Allocating currently reserved codes for vendor use.

3. A new TLV, EXTENDEDRESULT-TLV (0x118) that will carry a code



(which will be a superset of what is currently specified in [RFC 5812](#)) but also an optional cause content. This is illustrated in Figure 3.

#### [4.2.1.](#) New Codes

EXTENDEDRESULT-TLV Result Value is 32 bits and is a superset of [RFC 5810](#) Result TLV Result Value. The new version code space is 32 bits as opposed to the [RFC 5810](#) code size of 8 bits.

Code	Mnemonic	Details
0x100	E_EMPTY	Table is empty
0x101	E_INVALID_TFLAGS	Invalid table flags
0x102	E_INVALID_OP	Requested operation is invalid
0x103	E_CONGEST_NT	Node Congestion notification
0x104	E_COMPONENT_NOT_A_TABLE	Component not a table
0x105	E_PERM	Operation not permitted
0x107	E_BUSY	System is Busy
0x108	E_TIMED_OUT	A time out occurred while processing the message
0x106	E_UNKNOWN	A generic catch all error code. Carries a string to further extrapolate what the error implies.

Table 1: New codes

#### [4.2.2.](#) Vendor Codes

Codes 0x18-0xFE are reserved for use as vendor codes. Since these are freely available it is expected that the FE and CE side will both understand/interpret the semantics of any used codes.

#### [4.2.3.](#) Extended Result TLV

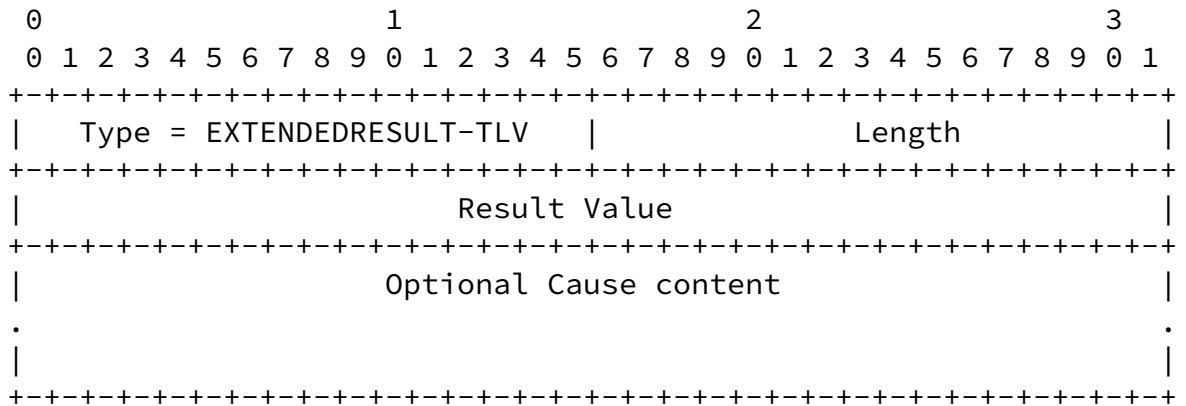


Figure 3: EXTENDEDRESULT-TLV

- o Like all other ForCES TLVs, the EXTENDEDRESULT-TLV is expected to be 32 bit aligned.
- o The Result Value derives and extends from the same current namespace as specified in [RFC 5810, section 7.1.7](#). The main difference is that we now have 32 bit result value (as opposed to the old 8 bit).
- o The optional result content is defined to further disambiguate the result value. It is expected Utf-8 values to be used. However, vendor specific error codes may choose to specify different contents. Additionally, future codes may specify cause contents to be of types other than string..
- o It is recommended that the maximum size of the cause string should not exceed 32 bytes. We do not propose the cause string be standardized.

XXX: Backward compatibility may require that we add a FEPO capability to advertise ability to do extended results so that the CE is able to interpret the results and a FEPO compatibility flag to define what TLV setting would be used. Alternatively, the backward compatibility can be made a configuration option (which helps reduce clutter on FEPO LFB given that it is expected that in the future it makes sense for implementations to support only EXTENDEDRESULT-TLVs).

### 4.3. Large Object Dumping

Imagine a GET request to a path that is a table i.e a table dump. Such a request is sent to the FE with a specific correlator, say X. Imagine this table to have a large number of entries at the FE. For the sake of illustration, lets say millions of rows. This requires

that the FE delivers the response over multiple messages, all using the same correlator X.

[RFC 5810](#) does not describe how a GET response is to indicate "I have more messages coming for this correlator".

Implementation experience indicates we can use the transaction flags to indicate that a GET response is the beginning, middle or end of a multi-part message. In other words we mirror the effect of an atomic transaction sent by a CE to an FE.

XXX: Add in the next update diagram and details of how this takes place.

## [5.](#) Acknowledgements

TBA

## [6.](#) IANA Considerations

This document registers two new top Level TLVs and two new path flags.

The following new TLVs are defined:

- o TABLERANGE-TLV (type ID 0x117)
- o EXTENDEDRESULT-TLV (type ID 0x118)

The following new path flags are defined:

- o F\_SELTABRANGE (value 0x2 i.e bit 1)

The Defined Result Values are changed:

- o codes 0x18-0xFE are reserved for vendor use.
- o codes 0x100-102 are defined by this document.

## 7. Security Considerations

TBD

## 8. References

Hadi Salim Expires September 18, 2014 [Page 10]

---

Internet-Draft ForCES Protocol Extensions March 2014

### 8.1. Normative References

- [RFC3746] Yang, L., Dantu, R., Anderson, T., and R. Gopal, "Forwarding and Control Element Separation (ForCES) Framework", [RFC 3746](#), April 2004.
- [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.

### 8.2. Informative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

### Author's Address

Jamal Hadi Salim  
Mojatatu Networks  
Suite 400, 303 Moodie Dr.  
Ottawa, Ontario K2H 9R4  
Canada

Email: [hadi@mojatatu.com](mailto:hadi@mojatatu.com)

Hadi Salim

Expires September 18, 2014

[Page 11]