

Workgroup: Network Working Group
Internet-Draft:
draft-dreibholz-rserpool-enrp-takeover-27
Published: 26 March 2023
Intended Status: Standards Track
Expires: 27 September 2023
Authors: T. Dreibholz X. Zhou
 SimulaMet Hainan University

Takeover Suggestion Flag for the ENRP Handle Update Message

Abstract

This document describes the Takeover Suggestion Flag for the ENRP_HANDLE_UPDATE message of the ENRP protocol.

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 <https://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 27 September 2023.

Copyright Notice

Copyright (c) 2023 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 (<https://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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

- [1. Introduction](#)
 - [1.1. Scope](#)
 - [1.2. Terminology](#)
 - [1.3. Conventions](#)
- [2. Takeover Suggestion Flag](#)
 - [2.1. Definition](#)
- [3. Reference Implementation](#)
- [4. Testbed Platform](#)
- [5. Security Considerations](#)
- [6. IANA Considerations](#)
- [7. References](#)
 - [7.1. Normative References](#)
 - [7.2. Informative References](#)
- [Authors' Addresses](#)

1. Introduction

Reliable Server Pooling as described in [3] defines protocols for providing highly available services. The management component used for pool administration is denoted as ENRP Server or Pool Registrar (PR). Since a single ENRP server constitutes a single point of failure, there must be multiple ENRP servers. Servers, denoted as Pool Elements (PE), use an arbitrary ENRP server for registration into the pool. The chosen ENRP server becomes the Home ENRP Server, also denoted as Home PR (PR-H), of the PE. It is responsible for making the PE identity known to the other ENRP servers (by using ENRP_HANDLE_UPDATE messages) and also to monitor the PE health (by using keep-alive messages).

As shown in [11], the following scenario leads to unbalanced ENRP server workload: consider a set of multiple ENRP servers with one subset being unreliable (for example, their network connection has problems) and some reliable ENRP servers. After a while, the reliable ENRP server will get the home ENRP server role for almost all of the PEs, which results in high workload for this ENRP server. Since the home ENRP server role is more computation-intensive (as shown by [13]), this leads to highly unbalanced workload for large RSerPool setups. This unbalanced workload remains, even when the unreliable ENRP servers become reliable again (for example, when the network problems have been solved).

1.1. Scope

The Takeover Suggestion Flag defined in this draft defines a flag for the ENRP_HANDLE_UPDATE message. If the flag is set, the receiving ENRP server is suggested to take over the PE specified in the ENRP_HANDLE_UPDATE message.

1.2. Terminology

The terms are commonly identified in related work and can be found in the RSerPool Overview document [RFC 5351](#) [3].

1.3. Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [[1](#)] [[2](#)] when, and only when, they appear in all capitals, as shown here.

2. Takeover Suggestion Flag

2.1. Definition

In this subsection, only the differences to the ENRP_HANDLE_UPDATE message defined in [5] are explained. The following figure shows the ENRP_HANDLE_UPDATE message:

[illegible]

T flag: 1 bit (boolean)

If set, the receiving ENRP server is suggested to take over the PE specified by the Pool Handle and Pool Element Parameters. It is RECOMMENDED for the receiving ENRP server to perform this takeover if it has the resources to do so.

3. Reference Implementation

The RSerPool reference implementation RSPLIB can be found at [15]. It supports the functionalities defined by [3], [4], [5], [6] and [8] as well as the options [9], [10] and of course the option

defined by this document. An introduction to this implementation is provided in [12].

4. Testbed Platform

A large-scale and realistic Internet testbed platform with support for the multi-homing feature of the underlying SCTP protocol is NorNet. A description of NorNet is provided in [14], some further information can be found on the project website [16].

5. Security Considerations

Security considerations for RSerPool systems are described by [7].

6. IANA Considerations

This document does not require additional IANA actions beyond those already identified in the ENRP and ASAP protocol specifications.

7. References

7.1. Normative References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [2] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [3] Lei, P., Ong, L., Tuexen, M., and T. Dreibholz, "An Overview of Reliable Server Pooling Protocols", RFC 5351, DOI 10.17487/RFC5351, September 2008, <<https://www.rfc-editor.org/info/rfc5351>>.
- [4] Stewart, R., Xie, Q., Stillman, M., and M. Tuexen, "Aggregate Server Access Protocol (ASAP)", RFC 5352, DOI 10.17487/RFC5352, September 2008, <<https://www.rfc-editor.org/info/rfc5352>>.
- [5] Xie, Q., Stewart, R., Stillman, M., Tuexen, M., and A. Silverton, "Endpoint Handlespace Redundancy Protocol (ENRP)", RFC 5353, DOI 10.17487/RFC5353, September 2008, <<https://www.rfc-editor.org/info/rfc5353>>.
- [6] Stewart, R., Xie, Q., Stillman, M., and M. Tuexen, "Aggregate Server Access Protocol (ASAP) and Endpoint Handlespace Redundancy Protocol (ENRP) Parameters", RFC

5354, DOI 10.17487/RFC5354, September 2008, <<https://www.rfc-editor.org/info/rfc5354>>.

- [7] Stillman, M., Ed., Gopal, R., Guttman, E., Sengodan, S., and M. Holdrege, "Threats Introduced by Reliable Server Pooling (RSerPool) and Requirements for Security in Response to Threats", RFC 5355, DOI 10.17487/RFC5355, September 2008, <<https://www.rfc-editor.org/info/rfc5355>>.
- [8] Dreibholz, T. and M. Tuexen, "Reliable Server Pooling Policies", RFC 5356, DOI 10.17487/RFC5356, September 2008, <<https://www.rfc-editor.org/info/rfc5356>>.
- [9] Dreibholz, T., "Handle Resolution Option for ASAP", Work in Progress, Internet-Draft, draft-dreibholz-rserpool-asap-hropt-29, 6 September 2021, <<https://www.ietf.org/archive/id/draft-dreibholz-rserpool-asap-hropt-29.txt>>.
- [10] Dreibholz, T. and X. Zhou, "Definition of a Delay Measurement Infrastructure and Delay-Sensitive Least-Used Policy for Reliable Server Pooling", Work in Progress, Internet-Draft, draft-dreibholz-rserpool-delay-28, 6 September 2021, <<https://www.ietf.org/archive/id/draft-dreibholz-rserpool-delay-28.txt>>.

7.2. Informative References

- [11] Zhou, X., Dreibholz, T., Fa, F., Du, W., and E. P. Rathgeb, "Evaluation and Optimization of the Registrar Redundancy Handling in Reliable Server Pooling Systems", Proceedings of the IEEE 23rd International Conference on Advanced Information Networking and Applications (AINA) Pages 256-262, ISBN 978-0-7695-3638-5, DOI 10.1109/AINA.2009.25, 26 May 2009, <<https://www.wiwi.uni-due.de/fileadmin/fileupload/I-TDR/ReliableServer/Publications/AINA2009.pdf>>.
- [12] Dreibholz, T., "Reliable Server Pooling – Evaluation, Optimization and Extension of a Novel IETF Architecture", 7 March 2007, <https://duepublico.uni-duisburg-essen.de/servlets/DerivateServlet/Derivate-16326/Dre2006_final.pdf>.
- [13] Dreibholz, T. and E. P. Rathgeb, "An Evaluation of the Pool Maintenance Overhead in Reliable Server Pooling Systems", SERSC International Journal on Hybrid Information Technology (IJHIT) Number 2, Volume 1, Pages 17-32, ISSN 1738-9968, April 2008, <<https://www.wiwi.uni->

due.de/fileadmin/fileupload/I-TDR/ReliableServer/Publications/IJHIT2008.pdf>.

- [14] Dreibholz, T. and E. G. Gran, "Design and Implementation of the NorNet Core Research Testbed for Multi-Homed Systems", Proceedings of the 3rd International Workshop on Protocols and Applications with Multi-Homing Support (PAMS) Pages 1094-1100, ISBN 978-0-7695-4952-1, DOI 10.1109/WAINA.2013.71, 27 March 2013, <<https://www.simula.no/file/threfereedinproceedingsreference2012-12-207643198512pdf/download>>.
- [15] Dreibholz, T., "Thomas Dreibholz's RSerPool Page", 2022, <<https://www.nntb.no/~dreibh/rserpool/>>.
- [16] Dreibholz, T., "NorNet – A Real-World, Large-Scale Multi-Homing Testbed", 2022, <<https://www.nntb.no/>>.

Authors' Addresses

Thomas Dreibholz
Simula Metropolitan Centre for Digital Engineering
Pilestredet 52
0167 Oslo
Norway

Email: dreibh@simula.no
URI: <https://www.simula.no/people/dreibh>

Xing Zhou
Hainan University, College of Information Science and Technology
Renmin Avenue 58
570228 Haikou
Hainan,
China

Phone: [+86-898-66279141](tel:+86-898-66279141)
Email: zhouxing@hainu.edu.cn