

Network Working Group
Internet-Draft
Intended status: Informational
Expires: 1 October 2025

T. Dreibholz
SimulaMet
M. Tuexen
Münster Univ. of App. Sciences
30 March 2025

Reliable Server Pooling (RSerPool) Bakeoff Scoring
draft-dreibholz-rserpool-score-36

Abstract

This memo describes some of the scoring to be used in the testing of Reliable Server Pooling protocols ASAP and ENRP at upcoming bakeoffs.

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 1 October 2025.

Copyright Notice

Copyright (c) 2025 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	2
2. Aggregate Server Access Protocol	2
2.1. Pool Element Communication	3
2.2. Pool User Communication	3
2.3. ENRP Server Communication	4
3. Endpoint Handlespace Redundancy Protocol	4
3.1. Peer Management	5
3.2. Update	5
3.3. Synchronization	6
3.4. Takeover	6
4. Bonus Points	7
5. Reference Implementation	7
6. Testbed Platform	7
7. Security Considerations	7
8. IANA Considerations	7
9. References	7
9.1. Normative References	7
9.2. Informative References	9
Authors' Addresses	9

1. Introduction

This document will be used as a basis for point scoring at upcoming RSerPool bakeoffs. Its purpose is similar to that described in RFC1025. It is hoped that a clear definition of where and how to score points will further the development of RSerPool.

Note that while attending a bakeoff no one else will score your points for you. We trust that all implementations will faithfully record their points that are received honestly. Note also that these scores are NOT to be used for marketing purposes. They are for the use of the implementations to know how well they are doing. The only reporting that will be done is a basic summary to the Reliable Server Pooling Working Group but please note that NO company or implementation names will be attached.

2. Aggregate Server Access Protocol

The ASAP protocol and useful extensions are described in the following documents:

* [3]

* [5]

* [9]

- * [10]

2.1. Pool Element Communication

These points will be scored for EACH peer implementation that you successfully communicate with.

- * 2 Successful ASAP Registration Request of a PE in a pool using Round Robin policy and handling of ASAP Registration Response.
- * 2 Failing ASAP Registration Request of a PE requesting Least Used policy in a pool using Round Robin policy and appropriate handling of ASAP Registration Response (e.g. printing error message, but not retrying registration).
- * 2 Successful re-registration of a PE in a pool using Round Robin policy.
- * 2 Successful ASAP Deregistration Request of the PE from its pool and handling of ASAP Deregistration Response.
- * 2 Successful handling of ASAP Endpoint Keep-Alive without Home bit set, i.e. answering with ASAP Endpoint Keep-Alive Ack.
- * 5 Successful handling of ASAP Endpoint Keep-Alive with Home bit set: respond with ASAP Endpoint Keep-Alive Ack and use new ENRP server for re-registration.
- * 5 Successful connection to and registration at an ENRP server announcing itself via multicast ASAP Announces.
- * 1 Successful registration into pool using Least Used policy.
- * 1 Successful registration into pool using Weighted Round Robin policy.
- * 1 Successful registration into pool using Random policy.
- * 1 Successful registration into pool using Weighted Random policy.

2.2. Pool User Communication

These points will be scored for EACH peer implementation that you successfully communicate with.

- * 5 Successful ASAP Handle Resolution in a pool using Round Robin policy, correct handling of ASAP Handle Resolution Response.

- * 2 Successful failure reporting using ASAP Endpoint Unreachable.
- * 5 Successful connection to and handle resolution at ENRP server announcing itself via multicast ASAP Announces.
- * 1 Successful handle resolution in a pool using Least Used policy.
- * 1 Successful handle resolution in a pool using Weighted Round Robin policy.
- * 1 Successful handle resolution in a pool using Random policy.
- * 1 Successful handle resolution in a pool using Weighted Random policy.

2.3. ENRP Server Communication

These points will be scored for EACH peer implementation that you successfully communicate with.

- * 2 Successful handling of an ASAP Registration Request into a pool using Round Robin policy (ENRP server answers with successful ASAP Registration Response).
- * 2 Rejecting registration of a PE requesting Round Robin policy into a pool using Least Used policy.
- * 5 Rejecting registration of a PE with all addresses **not** being part of the ASAP association.
- * 5 Successful registration of a PE with some addresses **not** being part of the ASAP association. The invalid addresses may **not** go into the handlespace.
- * 5 Successful handling of ASAP Endpoint Unreachable messages. The ENRP server must remove the given PE after MAX-BAD-PE-REPORTS=3 unreachability reports.
- * 2 Sending regular ASAP Endpoint Keep-Alives to its PEs.
- * 2 Removing PE not answering to ASAP Endpoint Keep-Alive.

3. Endpoint Handlespace Redundancy Protocol

The ENRP protocol and useful extensions are described in the following documents:

- * [4]

- * [5]

- * [11]

3.1. Peer Management

These points will be scored for EACH peer implementation that you successfully communicate with.

- * 2 Sending ENRP Presence to a new ENRP server.
- * 2 Sending ENRP Presences in the interval given by PEER-HEARTBEAT-CYCLE.
- * 5 Requesting peer list from new ENRP server using ENRP Peer List Request, handling ENRP Peer List Response and adding entries to its own peer list.
- * 2 Handling ENRP Peer List Request and replying with own peer list in ENRP Peer List Response.
- * 5 Requesting handlespace from new ENRP server using ENRP Handle Table Request, handling ENRP Handle Table Response (without M-bit set) and inserting entries into its own handlespace copy.
- * 5 Requesting handlespace from new ENRP server using ENRP Handle Table Request, handling ENRP Handle Table Response with M-bit set, requesting more entries and inserting entries into its own handlespace copy.
- * 2 Handling ENRP Handle Table Request and replying own handlespace in ENRP Handle Table Response (without M-bit).
- * 10 Handling ENRP Handle Table Request and replying own handlespace in ENRP Handle Table Response with M-bit set, remembering point to continue from, responding next block of handlespace entries upon following ENRP Handle Table Request, etc. until transfer of handlespace data is complete.
- * 5 Successful addition of new ENRP server announcing itself via multicast ENRP Presence (including association establishment as well as download of peer list and handlespace).

3.2. Update

These points will be scored for EACH peer implementation that you successfully communicate with.

- * 2 Handling an ENRP Handle Update adding a PE.
- * 2 Handling an ENRP Handle Update updating a PE. The changes must be entered into the local handlespace copy.
- * 2 Handling an ENRP Handle Update removing a PE.

3.3. Synchronization

These points will be scored for EACH peer implementation that you successfully communicate with.

- * 5 Successful detection of different handlespace checksums upon reception of ENRP Presence (due to additional PE), request of Handle Table with W-bit set, integration of missing PE into local handlespace copy and reporting the correct checksum in own ENRP Presence.
- * 5 Successful detection of different handlespace checksums upon reception of ENRP Presence (due to out-of-date PE), request of Handle Table with W-bit set, removal of PE from local handlespace copy and reporting the correct checksum in own ENRP Presence.
- * 10 Successful detection of different handlespace checksums upon reception of ENRP Presence (due to multiple new and out-of-date PE identities; size of PE identities is larger than maximum ENRP message size), request of Handle Table with W-bit set, handling of ENRP Handle Table Responses with M-bit set, removal of out-of-date PEs, integration of new PEs into the local handlespace copy and reporting correct checksum in own ENRP Presence.

3.4. Takeover

These points will be scored for EACH peer implementation that you successfully communicate with. The setup contains your ENRP server plus a set of peers running another implementation.

- * 5 Successfully detecting the failure of a remote peer and initiating a takeover procedure.
- * 5 Acknowledging another peer's takeover and aborting own takeover procedure.
- * 10 Correctly handling a remote peer's Takeover Server message, including ownership change for the remote peer's PEs.
- * 10 Successfully taking over a dead peer, including ownership change and informing the PEs taken over.

4. Bonus Points

You can also earn Bonus Points:

- * 20 points for the ENRP server handling the largest number of PEs.
- * 20 points for the ENRP server achieving the highest handle resolution throughput for a pool containing 100 (should this be larger?) PEs.

Please note that the whole period of the bakeoff is relevant.

5. Reference Implementation

The RSerPool reference implementation RSPLIB can be found at [14]. It supports the functionalities defined by [2], [3], [4], [5] and [7] as well as the options [9], [11] and [10]. The MIB module is defined in [8]. An introduction to this implementation is provided in [12].

6. 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 [13], some further information can be found on the project website [15].

7. Security Considerations

This document does only describe test scenarios and therefore does not introduce any new security issues.

For security considerations of the RSerPool protocols see [1], [2], [3], [4], [5], [7] and in particular [6].

8. IANA Considerations

This document introduces no additional considerations for IANA.

9. References

9.1. Normative References

- [1] Tuexen, M., Xie, Q., Stewart, R., Shore, M., Ong, L., Loughney, J., and M. Stillman, "Requirements for Reliable Server Pooling", RFC 3237, DOI 10.17487/RFC3237, January 2002, <<https://www.rfc-editor.org/info/rfc3237>>.

- [2] 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>>.
- [3] 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>>.
- [4] 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>>.
- [5] 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>>.
- [6] 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>>.
- [7] Dreibholz, T. and M. Tuexen, "Reliable Server Pooling Policies", RFC 5356, DOI 10.17487/RFC5356, September 2008, <<https://www.rfc-editor.org/info/rfc5356>>.
- [8] Dreibholz, T. and J. Mulik, "Reliable Server Pooling MIB Module Definition", RFC 5525, DOI 10.17487/RFC5525, April 2009, <<https://www.rfc-editor.org/info/rfc5525>>.
- [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>>.

- [11] Dreibholz, T. and X. Zhou, "Takeover Suggestion Flag for the ENRP Handle Update Message", Work in Progress, Internet-Draft, draft-dreibholz-rserpool-enrp-takeover-26, 6 September 2021, <<https://www.ietf.org/archive/id/draft-dreibholz-rserpool-enrp-takeover-26.txt>>.

9.2. Informative References

- [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. 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/threfreedinproceedingsreference2012-12-207643198512pdf/download>>.
- [14] Dreibholz, T., "Thomas Dreibholz's RSerPool Page", 2022, <<https://www.nntb.no/~dreibh/rserpool/>>.
- [15] 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
Stensberggata 27
0170 Oslo
Norway
Email: dreibh@simula.no
URI: <https://www.simula.no/people/dreibh>

Michael Tuexen
Mnster University of Applied Sciences
Stegerwaldstrae 39
48565 Steinfurt
Germany
Email: tuexen@fh-muenster.de