

Network Working Group
Internet-Draft
Intended status: Informational
Expires: 16 September 2026

T. Dreibholz
SimulaMet
15 March 2026

Ideas for a Next Generation of the Reliable Server Pooling Framework
draft-dreibholz-rserpool-nextgen-ideas-25

Abstract

This document collects some idea for a next generation of the Reliable Server Pooling framework.

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 16 September 2026.

Copyright Notice

Copyright (c) 2026 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
---------------------------	---

1.1. Abbreviations	2
1.2. Reliable Server Pooling	2
1.3. Scope	3
2. What to Change in the Next Generation of RSerPool?	3
2.1. Security Considerations	4
2.2. IANA Considerations	4
3. Reference Implementation	4
4. Testbed Platform	4
5. Acknowledgments	4
6. References	4
6.1. Normative References	4
6.2. Informative References	6
Author's Address	7

1. Introduction

1.1. Abbreviations

- * ASAP: Aggregate Server Access Protocol
- * ENRP: Endpoint Handlespace Redundancy Protocol
- * ID: Identifier
- * MPTCP: Multi-Path Transmission Control Protocol
- * PE: Pool Element
- * PR: Pool Registrar
- * PU: Pool User
- * RSerPool: Reliable Server Pooling
- * SCTP: Stream Control Transmission Protocol
- * TCP: Transmission Control Protocol
- * VNFPOOL: Virtual Network Function Resource Pooling

1.2. Reliable Server Pooling

Reliable Server Pooling (RSerPool) has been defined as RFCs in [1], [2], [3], [4], [5], [6], [7], [8]. There is also a detailed introduction provided by [16] as well as lots of further information material on [22]. RSerPool is therefore not introduced in more detail here.

1.3. Scope

The scope of this document is to collect some ideas of what to update/change for a next generation of the RSerPool framework. It is a result of lessons learned with more than one decade of RSerPool deployment (see also [17], [18], [19]) as well as ongoing discussions on applying RSerPool for Virtual Network Function Resource Pooling (VNFPOOL; as introduced in more detail in [15]).

2. What to Change in the Next Generation of RSerPool?

- * ENRP servers denote the management systems in the context of RSerPool. The term "server" is misleading, since ENRP servers are actually ENRP peers. Literature on RSerPool -- for example [16] -- therefore uses the more accurate term "Pool Registrar" (PR). A future revision of RSerPool should also use this term. (The RSerPool documents did not use "registrar" to avoid confusion with SIP registrars.)
- * Pool Element Identifiers (PE ID) and Pool Registrar Identifiers (PR ID) are 32-bit random numbers used for the identification of PEs and PRs. Since 64-bit CPUs are standard since quite a long time, these IDs should be extended to 64 bits.
- * ENRP uses the Internet-16 checksum defined in [9] to detect handlespace inconsistencies. It is trivially possible to extend the underlying algorithm to 32 bits, and the computation is more efficient on today's CPUs. The checksum algorithm should therefore be changed. (The Internet-16 checksum was finally chosen in 2005 after long decisions to avoid any possible patent issues. The trivial extension of Internet-16 to Internet-32 is probably not an issue any more?)
- * PR failures lead to possible concentration of all PUs and PEs at a single PR. To achieve a better load balancing, the solution ENRP Takeover Suggestion -- as defined in [14] -- should be included in ENRP.
- * For a Handle Resolution, a PR has to decide on how many PEs to select. Selecting too many ones causes additional overhead (which might be unnecessary), selecting too few ones may cause problems for the applications. The extension Handle Resolution Option -- defined in [12] -- provides a possibility for the PU to specify the amount of PEs to be selected. This possibility should be integrated into ASAP.

- * RSerPool defines SCTP (defined in [10]) as transport protocol for RSerPool. TCP and particularly Multi-Path TCP (MPTCP; see [11]) should be possible further transport protocols for all RSerPool traffic. SCTP should still be the recommended choice, but allowing TCP/MPTCP could make the deployment much easier. (SCTP is superior, but it lacks of support in operating systems and support by underlying network infrastructure, like firewalls and middleboxes.)

2.1. Security Considerations

Security considerations for RSerPool can be found in [6].

2.2. IANA Considerations

This document introduces no additional considerations for IANA.

3. Reference Implementation

The RSerPool reference implementation RSPLIB, including example applications, can be found at [22]. It supports the functionalities defined by [2], [3], [4], [5] and [6] as well as the options [12], [14] and [13]. An introduction to this implementation is provided in [16].

4. Testbed Platform

NorNet is a large-scale and realistic Internet testbed platform with support for Reliable Server Pooling as well as the underlying transport protocols SCTP and MPTCP. A description of and introduction to NorNet is provided in [19], [20], [21]. Further information can be found on the project website [23].

5. Acknowledgments

The author would like to thank Randall R. Stewart, Michael Tuexen, Ning Zong for their discussions and support.

6. References

6.1. 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] Braden, R., Borman, D., and C. Partridge, "Computing the Internet checksum", RFC 1071, DOI 10.17487/RFC1071, September 1988, <<https://www.rfc-editor.org/info/rfc1071>>.
- [10] Stewart, R., Ed., "Stream Control Transmission Protocol", RFC 4960, DOI 10.17487/RFC4960, September 2007, <<https://www.rfc-editor.org/info/rfc4960>>.
- [11] Ford, A., Raiciu, C., Handley, M., and O. Bonaventure, "TCP Extensions for Multipath Operation with Multiple Addresses", RFC 6824, DOI 10.17487/RFC6824, January 2013, <<https://www.rfc-editor.org/info/rfc6824>>.

- [12] 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>>.
- [13] 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>>.
- [14] 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>>.
- [15] Zong, N., Dunbar, L., Shore, M., Lopez, D., and G. Karagiannis, "Virtualized Network Function (VNF) Pool Problem Statement", Work in Progress, Internet-Draft, draft-zong-vnfpool-problem-statement-06, 1 July 2014, <<https://www.ietf.org/archive/id/draft-zong-vnfpool-problem-statement-06.txt>>.

6.2. Informative References

- [16] 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>.
- [17] Dreibholz, T. and E. P. Rathgeb, "A Powerful Tool-Chain for Setup, Distributed Processing, Analysis and Debugging of OMNeT++ Simulations", Proceedings of the 1st ACM/ICST International Workshop on OMNeT++ ISBN 978-963-9799-20-2, DOI 10.4108/ICST.SIMUTOOLS2008.2990, 7 March 2008, <<https://www.wiwi.uni-due.de/fileadmin/fileupload/I-TDR/ReliableServer/Publications/OMNeTWorkshop2008.pdf>>.
- [18] Dreibholz, T., "Evaluation and Optimisation of Multi-Path Transport using the Stream Control Transmission Protocol", Habilitation Treatise, 13 March 2012, <https://duepublico.uni-duisburg-essen.de/servlets/DerivateServlet/Derivate-29737/Dre2012_final.pdf>.

- [19] 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>>.
- [20] Dreibholz, T., "The NorNet Core Testbed Introduction and Status", Proceedings of the 1st International NorNet Users Workshop (NNUW-1), 18 September 2013, <<https://www.simula.no/file/simulasimula2124pdf/download>>.
- [21] Dreibholz, T., "The NorNet Core Testbed An Experiment Tutorial", Proceedings of the 1st International NorNet Users Workshop (NNUW-1), 19 September 2013, <<https://www.simula.no/file/simulasimula2130pdf/download>>.
- [22] Dreibholz, T., "Thomas Dreibholz's RSerPool Page", 2022, <<https://www.nntb.no/~dreibh/rserpool/>>.
- [23] Dreibholz, T., "NorNet A Real-World, Large-Scale Multi-Homing Testbed", 2022, <<https://www.nntb.no/>>.

Author's Address

Thomas Dreibholz
Simula Metropolitan Centre for Digital Engineering
Stensberggata 27
0170 Oslo
Norway
Email: dreibh@simula.no
URI: <https://www.simula.no/people/dreibh>