

Network Working Group  
Internet-Draft  
Expires: Januar 4, 2005

M. Tuexen  
Univ. of Applied Sciences Muenster  
T. Dreibholz  
University of Duisburg-Essen  
July 6, 2004

**Reliable Server Pooling Policies**  
**draft-tuexen-rserpool-policies-00.txt**

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>.

This Internet-Draft will expire on Januar 4, 2005.

Copyright Notice

Copyright (C) The Internet Society (2004). All Rights Reserved.

Abstract

This document describes server pool policies for Reliable Server Pooling including considerations for implementing them at name servers and pool users.

## Table of Contents

<a href="#">1.</a>	<a href="#">Introduction . . . . .</a>	<a href="#">4</a>
<a href="#">2.</a>	<a href="#">Terminology and Definitions . . . . .</a>	<a href="#">4</a>
<a href="#">2.1</a>	<a href="#">Load . . . . .</a>	<a href="#">4</a>
<a href="#">2.2</a>	<a href="#">Weight . . . . .</a>	<a href="#">4</a>
<a href="#">3.</a>	<a href="#">Static Policies . . . . .</a>	<a href="#">5</a>
<a href="#">3.1</a>	<a href="#">Round Robin Policy . . . . .</a>	<a href="#">5</a>
<a href="#">3.1.1</a>	<a href="#">Description . . . . .</a>	<a href="#">5</a>
<a href="#">3.1.2</a>	<a href="#">Name Server Considerations . . . . .</a>	<a href="#">5</a>
<a href="#">3.1.3</a>	<a href="#">Pool User Considerations . . . . .</a>	<a href="#">5</a>
<a href="#">3.1.4</a>	<a href="#">Pool Member Selection Policy Parameter . . . . .</a>	<a href="#">5</a>
<a href="#">3.2</a>	<a href="#">Weighted Round Robin Policy . . . . .</a>	<a href="#">5</a>
<a href="#">3.2.1</a>	<a href="#">Description . . . . .</a>	<a href="#">5</a>
<a href="#">3.2.2</a>	<a href="#">Name Server Considerations . . . . .</a>	<a href="#">6</a>
<a href="#">3.2.3</a>	<a href="#">Pool User Considerations . . . . .</a>	<a href="#">6</a>
<a href="#">3.2.4</a>	<a href="#">Pool Member Selection Policy Parameter . . . . .</a>	<a href="#">6</a>
<a href="#">3.3</a>	<a href="#">Random Policy . . . . .</a>	<a href="#">6</a>
<a href="#">3.3.1</a>	<a href="#">Description . . . . .</a>	<a href="#">6</a>
<a href="#">3.3.2</a>	<a href="#">Name Server Considerations . . . . .</a>	<a href="#">6</a>
<a href="#">3.3.3</a>	<a href="#">Pool User Considerations . . . . .</a>	<a href="#">6</a>
<a href="#">3.3.4</a>	<a href="#">Pool Member Selection Policy Parameter . . . . .</a>	<a href="#">7</a>
<a href="#">3.4</a>	<a href="#">Weighted Random Policy . . . . .</a>	<a href="#">7</a>
<a href="#">3.4.1</a>	<a href="#">Description . . . . .</a>	<a href="#">7</a>
<a href="#">3.4.2</a>	<a href="#">Name Server Considerations . . . . .</a>	<a href="#">7</a>
<a href="#">3.4.3</a>	<a href="#">Pool User Considerations . . . . .</a>	<a href="#">7</a>
<a href="#">3.4.4</a>	<a href="#">Pool Member Selection Policy Parameter . . . . .</a>	<a href="#">7</a>
<a href="#">4.</a>	<a href="#">Dynamic Policies . . . . .</a>	<a href="#">8</a>
<a href="#">4.1</a>	<a href="#">Least Used Policy . . . . .</a>	<a href="#">8</a>
<a href="#">4.1.1</a>	<a href="#">Description . . . . .</a>	<a href="#">8</a>
<a href="#">4.1.2</a>	<a href="#">Name Server Considerations . . . . .</a>	<a href="#">8</a>
<a href="#">4.1.3</a>	<a href="#">Pool User Considerations . . . . .</a>	<a href="#">8</a>
<a href="#">4.1.4</a>	<a href="#">Pool Member Selection Policy Parameter . . . . .</a>	<a href="#">8</a>
<a href="#">4.2</a>	<a href="#">Least Used with Degradation Policy . . . . .</a>	<a href="#">8</a>
<a href="#">4.2.1</a>	<a href="#">Description . . . . .</a>	<a href="#">8</a>
<a href="#">4.2.2</a>	<a href="#">Name Server Considerations . . . . .</a>	<a href="#">9</a>
<a href="#">4.2.3</a>	<a href="#">Pool User Considerations . . . . .</a>	<a href="#">9</a>
<a href="#">4.2.4</a>	<a href="#">Pool Member Selection Policy Parameter . . . . .</a>	<a href="#">9</a>
<a href="#">4.3</a>	<a href="#">Priority Least Used Policy . . . . .</a>	<a href="#">9</a>
<a href="#">4.3.1</a>	<a href="#">Description . . . . .</a>	<a href="#">9</a>
<a href="#">4.3.2</a>	<a href="#">Name Server Considerations . . . . .</a>	<a href="#">10</a>
<a href="#">4.3.3</a>	<a href="#">Pool User Considerations . . . . .</a>	<a href="#">10</a>
<a href="#">4.3.4</a>	<a href="#">Pool Member Selection Policy Parameter . . . . .</a>	<a href="#">10</a>
<a href="#">4.4</a>	<a href="#">Randomized Least Used Policy . . . . .</a>	<a href="#">10</a>
<a href="#">4.4.1</a>	<a href="#">Description . . . . .</a>	<a href="#">10</a>
<a href="#">4.4.2</a>	<a href="#">Name Server Considerations . . . . .</a>	<a href="#">10</a>
<a href="#">4.4.3</a>	<a href="#">Pool User Considerations . . . . .</a>	<a href="#">11</a>
<a href="#">4.4.4</a>	<a href="#">Pool Member Selection Policy Parameter . . . . .</a>	<a href="#">11</a>



<a href="#">5.</a>	Security Considerations . . . . .	<a href="#">11</a>
<a href="#">6.</a>	IANA Considerations . . . . .	<a href="#">11</a>
<a href="#">7.</a>	References . . . . .	<a href="#">11</a>
<a href="#">7.1</a>	Normative References . . . . .	<a href="#">11</a>
<a href="#">7.2</a>	Informative References . . . . .	<a href="#">12</a>
	Authors' Addresses . . . . .	<a href="#">12</a>
	Intellectual Property and Copyright Statements . . . . .	<a href="#">13</a>

## **1. Introduction**

The protocols defined in ENRP [5], ASAP [4] and Parameters [3] support a variety of server policies. Some of the policies use dynamic load information of the pool elements and others do not. Therefore, we classify them as dynamic and static. The selection of the pool user is performed by two different entities. Some of the consequences for policies which are not stateless are described in Performance [7].

Therefore this document describes not only packet formats but also gives a detailed description of the procedures to be followed at the name servers and the pool users to implement each server policy.

## **2. Terminology and Definitions**

### **2.1 Load**

The term load is a value specifying how much a pool element's resources are currently utilized. 0x000000 states, that the pool element is not utilized (0%), 0xffffffff states that it is fully utilized (100%). Defining what utilization means is application-dependent and out of the scope of RSerPool. However, it is required that all pool elements of the same pool using load information have the same definition of load.

For example, load may define the current amount of users out of a maximum on a FTP server, the CPU usage of a database server or the memory utilization of a compute service.

### **2.2 Weight**

Weight defines a pool element's service capacity relatively to other pool elements of the same pool. Theoretically, there is no upper limit for weight values (although limited by datatype size). Defining what value weights compare is application-dependent and out of the scope of RSerPool. However, it is required that all pool elements of the same pool using weight information have the same definition of weight.

A weight of 0 denotes that the pool element is not capable of providing any service, a weight of  $2^n$  denotes that the pool element is capable of providing a two times better service than a pool element having weight  $n$ .

For example, weight may define a compute service's computation capacity. That is, a pool element of weight 100 will complete a work package in half of the time compared to a pool element of weight 50.



### **3. Static Policies**

#### **3.1 Round Robin Policy**

##### **3.1.1 Description**

The Round Robin (RR) policy is a very simple and efficient policy which requires state. This policy is denoted as the default policy and MUST be supported by all RSerPool components.

##### **3.1.2 Name Server Considerations**

The name server SHOULD hold the pool elements of each server pool in a circular list and SHOULD store a pointer to one of the elements, called the head. On reception of a name resolution request the name server SHOULD return the pool elements from the circular list starting with head. Then head SHOULD be advanced by one element.

Using this algorithm it is made sure that not all lists presented to the pool users start with the same element.

##### **3.1.3 Pool User Considerations**

A pool user SHOULD use the list of pool elements returned by the name server in a round robin fashion, starting with the first. If all elements of the list have been used it should start from the beginning again until the information is out of date.

##### **3.1.4 Pool Member Selection Policy Parameter**

```

      0               1               2               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Param Type = 0x6               |   Length = 0x8               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Policy=0x1   |               (reserved)               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

- o Reserved: 24 bits, SHOULD be set to 0.

#### **3.2 Weighted Round Robin Policy**

##### **3.2.1 Description**

The Weighted Round Robin (WRR) policy is a generalization of the RR policy. If all weights are 1 then WRR is just RR.



### [3.2.2](#) Name Server Considerations

The name server SHOULD follow the same rules as for RR but initialize and modify the circular list differently. The name server puts each pool element possibly multiple times into the list such that:

- o The ratio of the number of occurrences of a pool element to the list length is the same as the ratio of the weight of that pool element to the sum of weights.
- o Each pool element is inserted as distributed as possible in the circular list.

### [3.2.3](#) Pool User Considerations

The pool user SHOULD follow the same rules as for RR.

### [3.2.4](#) Pool Member Selection Policy Parameter

```

      0              1              2              3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|   Param Type = 0x6               |   Length = 0x8               |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Policy=0x2   |                   Weight                         |
+--+--+--+--+--+--+--+--+--+--+--+-----+

```

- o Weight: Weight constant for the WRR process.

## [3.3](#) Random Policy

### [3.3.1](#) Description

The Random (RAND) policy is a very simple stateless policy.

### [3.3.2](#) Name Server Considerations

The name server selects at most the requested number of pool elements from the list of pool elements. Each element MUST NOT be reported more than once to the pool user.

### [3.3.3](#) Pool User Considerations

Each time the pool user must select one pool element it does this by randomly selecting one element from the list of pool elements received from the name server.



### [3.3.4](#) Pool Member Selection Policy Parameter

```

      0                               1                               2                               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Param Type = 0x6                               |   Length = 0x8   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Policy=0x3   |                               (reserved)           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

- o Reserved: 24 bits, SHOULD be set to 0.

## [3.4](#) Weighted Random Policy

### [3.4.1](#) Description

The Weighted Random (WRAND) policy is a generalization of the RAND policy, adding a weight for each pool element entry. RAND is equal to WRAND having all weights set to 1.

### [3.4.2](#) Name Server Considerations

The name server SHOULD select at most the requested number of pool elements randomly from the list of pool elements. Each element MUST NOT be reported more than once to the pool user. The probability of selecting a pool element should be the ratio of the weight of that pool element to the sum of weights.

### [3.4.3](#) Pool User Considerations

Each time the pool user must select one pool element it does this by randomly selecting one element from the list of pool elements received from the name server.

### [3.4.4](#) Pool Member Selection Policy Parameter

```

      0                               1                               2                               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Param Type = 0x6                               |   Length = 0x8   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Policy=0x4   |                               Weight                |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

- o Weight: Weight constant for the WRAND process.



## **4. Dynamic Policies**

### **4.1 Least Used Policy**

#### **4.1.1 Description**

The Least Used (LU) policy uses load information provided by the pool elements to select the lowest-loaded pool elements within the pool.

#### **4.1.2 Name Server Considerations**

The name server SHOULD select at most the requested number of pool elements. Their load values SHOULD be the lowest possible ones within the pool. Each element MUST NOT be reported more than once to the pool user. If there is a choice of equal-loaded pool elements, round robin selection SHOULD be made between these elements. The returned list of pool elements MUST be sorted ascending by load value.

#### **4.1.3 Pool User Considerations**

The pool user should try to use the pool elements returned from the list in the order returned by the name server. A subsequent call for name resolution may result in the same list. Therefore, it is RECOMMENDED for a pool user to request multiple entries in order to have a sufficient amount of feasible backup entries available.

#### **4.1.4 Pool Member Selection Policy Parameter**

```

      0              1              2              3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|   Param Type = 0x6              |   Length = 0x8              |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Policy=0x5   |              Load              |
+--+--+--+--+--+--+--+--+--+--+--+-----+

```

- o Load: Current load of the pool element.

### **4.2 Least Used with Degradation Policy**

#### **4.2.1 Description**

The Least Used with Degradation (LUD) policy extends the LU policy by a load degradation value describing the pool element's load increment when a new service association is accepted.



Example: Pool elements A and B are loaded by 50%, but the load of A will increase due to a new application request only by 10% while B will be fully loaded. PLU allows to specify this load degradation in the policy information, the selection is made on the lowest sum of load and degradation value. That is, A will be selected ( $50+10=60$ ) instead of B ( $50+50=100$ ).



### 4.3.2 Name Server Considerations

The name server SHOULD select at most the requested number of pool elements. Their sums of load + degradation SHOULD be the lowest possible ones within the pool. Each element MUST NOT be reported more than once to the pool user. If there is a choice of equal-valued pool element entries, round robin SHOULD be made between these elements. The returned list of pool elements MUST be sorted ascending by the sum of load and degradation value.

### 4.3.3 Pool User Considerations

The pool user should try to use the pool elements returned from the list in the order returned by the name server. A subsequent call for name resolution may result in the same list. Therefore, it is RECOMMENDED for a pool user to request multiple entries in order to have a sufficient amount of feasible backup entries available.

#### 4.3.4 Pool Member Selection Policy Parameter

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+									
Param Type = 0x6										Length = 0xc																													
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+									
Policy=0x7										Load																													
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+									
(reserved)										Load Degradation																													
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+									

- o Load: Current load of the pool element.
- o Load Degradation: Load Degradation constant of the pool element.

#### 4.4 Randomized Least Used Policy

#### 4.4.1 Description

The Randomized Least Used (RLU) policy combines LU and WRAND. That is, the pool element entries are selected randomly; the probability for a pool element entry to be selected is the ratio of 100%-load to the sum of all pool elements' load values.

#### 4.4.2 Name Server Considerations

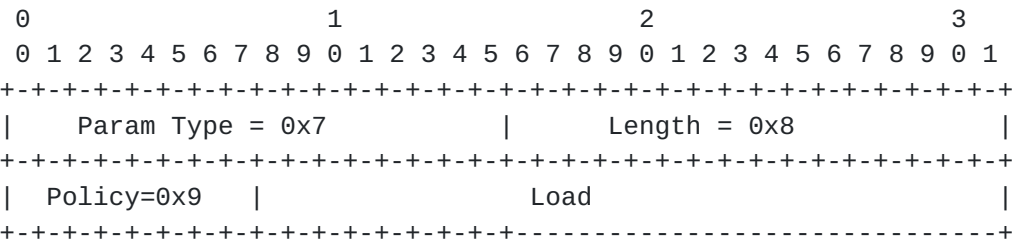
The name server SHOULD behave like WRAND, having every PE's weight set to (0xffffffff - Load value provided by the pool element).



4.4.3 Pool User Considerations

See WRAND policy.

4.4.4 Pool Member Selection Policy Parameter



o Load: Current load of the pool element.

5. Security Considerations

The security threats regarding RSerPool have been analyzed in RSerPool threats [6]. The server policy descriptions in this document do not add any other threats.

6. IANA Considerations

IANA keeps a list of Policy Types which are 1 byte values. The Policy values used in this document are:

Value	Policy
-----	-----
0x00	(reserved by IETF)
0x01	Round Robin
0x02	Weighted Round Robin
0x03	Random
0x04	Weighted Random
0x05	Least Used
0x06	Least Used with Degradation
0x07	Priority Least Used
0x09	Randomized Least Used
others	(reserved by IETF)

7. References

7.1 Normative References

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



- [2] Bradner, S., "Intellectual Property Rights in IETF Technology", [BCP 79](#), [RFC 3668](#), February 2004.
- [3] Stewart, R., Xie, Q. and M. Tuexen, "Aggregate Server Access Protocol (ASAP) and Endpoint Name Resolution (ENRP) Parameters", [draft-ietf-rserpool-common-param-05](#) (work in progress), October 2003.
- [4] Stewart, R., Xie, Q., Stillman, M. and M. Tuexen, "Aggregate Server Access Protocol (ASAP)", [draft-ietf-rserpool-asap-08](#) (work in progress), October 2003.
- [5] Xie, Q., Stewart, R. and M. Stillman, "Endpoint Name Resolution Protocol (ENRP)", [draft-ietf-rserpool-enrp-07](#) (work in progress), October 2003.
- [6] Stillman, M., "Threats Introduced by Rserpool and Requirements for Security in response to Threats", [draft-ietf-rserpool-threats-02](#) (work in progress), October 2003.

## **[7.2](#) Informative References**

- [7] Dreibholz, T., Rathgeb, E. and M. Tuexen, "Load Distribution Performance of the Reliable Server Pooling Framework", Submitted to Globecom 2004.

### Authors' Addresses

Michael Tuexen  
University of Applied Sciences Muenster  
Stegerwaldstrasse 39  
48565 Steinfurt, Nordrhein-Westfalen  
Germany

EMail: [tuexen@fh-muenster.de](mailto:tuexen@fh-muenster.de)

Thomas Dreibholz  
University of Duisburg-Essen, Institute for Experimental Mathematics  
Ellernstrasse 29  
45326 Essen, Nordrhein-Westfalen  
Germany

Phone: +49 201 183-7637

Fax: +49 201 183-7673

EMail: [dreibh@exp-math.uni-essen.de](mailto:dreibh@exp-math.uni-essen.de)

URI: <http://www.exp-math.uni-essen.de/~dreibh/>



## Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the IETF's procedures with respect to rights in IETF Documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at [ietf-ipr@ietf.org](mailto:ietf-ipr@ietf.org).

## Disclaimer of Validity

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

## Copyright Statement

Copyright (C) The Internet Society (2004). This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

## Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.

