

Network Working Group  
Request for Comments: 5355  
Category: Informational

M. Stillman, Ed.  
Nokia  
R. Gopal  
Nokia Siemens Networks  
E. Guttman  
Sun Microsystems  
S. Sengodan  
Nokia Siemens Networks  
M. Holdrege  
September 2008

Threats Introduced by Reliable Server Pooling (RSerPool)  
and Requirements for Security in Response to Threats

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Abstract

Reliable Server Pooling (RSerPool) is an architecture and set of protocols for the management and access to server pools supporting highly reliable applications and for client access mechanisms to a server pool. This document describes security threats to the RSerPool architecture and presents requirements for security to thwart these threats.

## Table of Contents

|  |    |
|--|----|
| 1. Introduction .....  | 3  |
| 1.1. Definitions .....   | 3  |
| 1.2. Conventions .....   | 4  |
| 2. Threats .....   | 4  |
| 2.1. PE Registration/De-Registration Flooding --<br>Non-Existent PE .....            | 4  |
| 2.2. PE Registration/De-Registration Flooding --<br>Unauthorized PE .....            | 5  |
| 2.3. PE Registration/De-Registration Spoofing .....                                  | 6  |
| 2.4. PE Registration/De-Registration Unauthorized .....                              | 6  |
| 2.5. Malicious ENRP Server Joins the Group of Legitimate<br>ENRP Servers .....       | 7  |
| 2.6. Registration/De-Registration with Malicious ENRP Server .....                   | 7  |
| 2.7. Malicious ENRP Handlespace Resolution .....                                     | 8  |
| 2.8. Malicious Node Performs a Replay Attack .....                                   | 9  |
| 2.9. Re-Establishing PU-PE Security during Failover .....                            | 9  |
| 2.10. Integrity .....  | 10 |
| 2.11. Data Confidentiality .....   | 10 |
| 2.12. ENRP Server Discovery .....  | 11 |
| 2.13. Flood of Endpoint-Unreachable Messages from the<br>PU to the ENRP Server ..... | 12 |
| 2.14. Flood of Endpoint Keep-Alive Messages from the<br>ENRP Server to a PE .....    | 12 |
| 2.15. Security of the ENRP Database .....  | 13 |
| 2.16. Cookie Mechanism Security .....  | 13 |
| 2.17. Potential Insider Attacks from Legitimate ENRP Servers .....                   | 14 |
| 3. Security Considerations .....   | 15 |
| 4. Normative References .....  | 17 |

## 1. Introduction

The RSerPool architecture [RFC5351] supports high-availability and load balancing by enabling a pool user to identify the most appropriate server from the server pool at a given time. The architecture is defined to support a set of basic goals. These include application-independent protocol mechanisms, separation of server naming from IP addressing, the use of the end-to-end principle to avoid dependencies on intermediate equipment, separation of session availability/failover functionality from the application itself, the ability to facilitate different server selection policies, the ability to facilitate a set of application-independent failover capabilities, and a peer-to-peer structure.

RSerPool provides a session layer for robustness. The session layer function may redirect communication transparently to upper layers. This alters the direct one-to-one association between communicating endpoints that typically exists between clients and servers. In particular, secure operation of protocols often relies on assumptions at different layers regarding the identity of the communicating party and the continuity of the communication between endpoints. Further, the operation of RSerPool itself has security implications and risks. The session layer operates dynamically, which imposes additional concerns for the overall security of the end-to-end application.

This document explores the security implications of RSerPool, both due to its own functions and due to its being interposed between applications and transport interfaces.

This document is related to the RSerPool Aggregate Server Access Protocol (ASAP) [RFC5352] and Endpoint Name Resolution Protocol (ENRP) [RFC5353] documents, which describe, in their Security Consideration sections, the mechanisms for meeting the security requirements in this document. TLS [RFC5246] is the security mechanism for RSerPool that was selected to meet all the requirements described in this document. The Security Considerations sections of ASAP and ENRP describe how TLS is actually used to provide the security that is discussed in this document.

### 1.1. Definitions

This document uses the following terms:

Endpoint Name Resolution Protocol (ENRP):

Within the operational scope of RSerPool, ENRP[RFC5353] defines the procedures and message formats of a distributed fault-tolerant registry service for storing, bookkeeping, retrieving, and distributing pool operation and membership information.

**Aggregate Server Access Protocol (ASAP):**

ASAP [RFC5352] is a session layer protocol that uses ENRP to provide a high-availability handlespace. ASAP is responsible for the abstraction of the underlying transport technologies, load distribution management, fault management, as well as the presentation to the upper layer (i.e., the ASAP User) of a unified primitive interface.

**Operational scope:**

The part of the network visible to pool users by a specific instance of the Reliable Server Pooling protocols.

**Pool (or server pool):**

A collection of servers providing the same application functionality.

**Pool handle:**

A logical pointer to a pool. Each server pool will be identifiable in the operational scope of the system by a unique pool handle.

**ENRP handlespace (or handlespace):**

A cohesive structure of pool names and relations that may be queried by a client. A client in this context is an application that accesses another remote application running on a server using a network.

**Pool element (PE):** A server entity having registered to a pool.

**Pool user (PU):** A server pool user.

## 1.2. Conventions

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].

## 2. Threats

### 2.1. PE Registration/De-Registration Flooding -- Non-Existent PE

#### 2.1.1. Threat

A malicious node could send a stream of false registrations/de-registrations on behalf of non-existent PEs to ENRP servers at a very rapid rate and thereby create unnecessary state in an ENRP server.

### 2.1.2. Effect

The malicious node will corrupt the pool registrar database and/or disable the RSerPool discovery and database function. This represents a denial-of-service attack, as the PU would potentially get an IP address of a non-existent PE in response to an ENRP query.

### 2.1.3. Requirement

An ENRP server that receives a registration/de-registration MUST NOT create or update state information until it has authenticated the PE. TLS with a pre-shared-key (PSK) is mandatory to implement as the authentication mechanism. For PSK, having a pre-shared-key constitutes authorization. The network administrators of a pool need to decide which nodes are authorized to participate in the pool. The justification for PSK is that we assume that one administrative domain will control and manage the server pool. This allows for PSK to be implemented and managed by a central security administrator.

## 2.2. PE Registration/De-Registration Flooding -- Unauthorized PE

### 2.2.1. Threat

A malicious node or PE could send a stream of registrations/de-registrations that are unauthorized to register/de-register to ENRP servers at a very rapid rate and thereby create unnecessary state in an ENRP server.

### 2.2.2. Effect

This attack will corrupt the pool registrar database and/or disable the RSerPool discovery and database function. There is the potential for two types of attacks: denial of service and data interception. In the denial-of-service attack, the PU gets an IP address of a rogue PE in response to an ENRP query, which might not provide the actual service. In addition, a flood of message could prevent legitimate PEs from registering. In the data interception attack, the rogue PE does provide the service as a man in the middle (MITM), which allows the attacker to collect data.

### 2.2.3. Requirement

An ENRP server that receives a registration/de-registration MUST NOT create or update state information until the authentication information of the registering/de-registering entity is verified.

TLS is used as the authentication mechanism between the ENRP server and PE. TLS with PSK is mandatory to implement as the authentication mechanism. For PSK, having a pre-shared-key constitutes authorization. The network administrators of a pool need to decide which nodes are authorized to participate in the pool.

### 2.3. PE Registration/De-Registration Spoofing

#### 2.3.1. Threat

A malicious node could send false registrations/de-registrations to ENRP servers concerning a legitimate PE, thereby creating false state information in the ENRP servers.

#### 2.3.2. Effect

This would generate misinformation in the ENRP server concerning a PE and would be propagated to other ENRP servers, thereby corrupting the ENRP database. Distributed Denial of Service (DDoS) could result: if a PE that is a target for a DDoS attack for some popular high-volume service, then the attacker can register a PE to which a lot of PUs will try to connect. This allows man-in-the-middle or masquerade attacks on the service provided by the legitimate PEs. If an attacker registers its server address as a PE and handles the requests, he can eavesdrop on service data.

#### 2.3.3. Requirement

An ENRP server that receives a registration/de-registration MUST NOT create or update state information until it has authenticated the PE. TLS is used as the authentication mechanism between the ENRP server and the PE. TLS with PSK is mandatory to implement as the authentication mechanism. For PSK, having a pre-shared-key constitutes authorization. The network administrators of a pool need to decide which nodes are authorized to participate in the pool. A PE can register only for itself and cannot register on behalf of other PEs.

### 2.4. PE Registration/De-Registration Unauthorized

#### 2.4.1. Threat

A PE that is not authorized to join a pool could send registrations/de-registrations to ENRP servers, thereby creating false state information in the ENRP servers.

#### 2.4.2. Effect

This attack would generate misinformation in the ENRP server concerning a PE and would be propagated to other ENRP servers thereby corrupting the ENRP database. This allows man-in-the-middle or masquerade attacks on the service provided by the legitimate PEs. If an attacker registers its server address as a PE and handles the requests, he can eavesdrop on service data.

#### 2.4.3. Requirement

An ENRP server that receives a registration/de-registration MUST NOT create or update state information until it has authorized the requesting entity. TLS is used as the authentication mechanism. TLS with PSK is mandatory to implement as the authentication mechanism. For PSK, having a pre-shared-key constitutes authorization. The network administrators of a pool need to decide which nodes are authorized to participate in the pool.

### 2.5. Malicious ENRP Server Joins the Group of Legitimate ENRP Servers

#### 2.5.1. Threat

A malicious ENRP server joins the group of legitimate ENRP servers with the intent of propagating inaccurate updates to corrupt the ENRP database. The attacker sets up an ENRP server and attempts to communicate with other ENRP servers.

#### 2.5.2. Effect

The result would be Inconsistent ENRP database state.

#### 2.5.3. Requirement

ENRP servers MUST perform mutual authentication. This would prevent the attacker from joining its ENRP server to the pool. TLS is used as the mutual authentication mechanism. TLS with PSK is mandatory to implement as the authentication mechanism. For PSK, having a pre-shared-key constitutes authorization. The network administrators of a pool need to decide which nodes are authorized to participate in the pool.

### 2.6. Registration/De-Registration with Malicious ENRP Server

#### 2.6.1. Threat

A PE unknowingly registers/de-registers with a malicious ENRP server.

### 2.6.2. Effect

The registration might not be properly processed or it might be ignored. A rogue ENRP server has the ability to return any address to a user requesting service; this ability could result in denial of service or connection to a rogue PE that is the attacker's choice for service.

### 2.6.3. Requirement

The PE MUST authenticate the ENRP server. TLS is the mechanism used for the authentication. TLS with PSK is mandatory to implement as the authentication mechanism. For PSK, having a pre-shared-key constitutes authorization. The network administrators of a pool need to decide which nodes are authorized to participate in the pool. This requirement prevents malicious outsiders and insiders from adding their own ENRP server to the pool.

## 2.7. Malicious ENRP Handlespace Resolution

### 2.7.1. Threat

The ASAP protocol receives a handlespace resolution response from an ENRP server, but the ENRP server is malicious and returns random IP addresses or an inaccurate list in response to the pool handle.

### 2.7.2. Effect

The PU application communicates with the wrong PE or is unable to locate the PE since the response is incorrect in saying that a PE with that handle did not exist. A rogue ENRP server has the ability to return any address to ASAP requesting an address list that could result in denial of service or connection to a rogue PE of the attacker's choice for service. From the PE, the attacker could eavesdrop or tamper with the application.

### 2.7.3. Requirement

ASAP SHOULD authenticate the ENRP server. TLS with certificates is the mandatory-to-implement mechanism used for authentication. The administrator uses a centralized Certificate Authority (CA) to generate and sign certificates. The certificate is stored on the ENRP server. A CA trusted root certification authority certificate is sent to the client out of band, and the certificate signature on the ENRP server certificate is checked for validity during the TLS handshake. This authentication prevents malicious outsiders and insiders from adding an ENRP server to the pool that may be accessed by ASAP.



## 2.8. Malicious Node Performs a Replay Attack

### 2.8.1. Threat

A malicious node could replay the entire message previously sent by a legitimate entity. This could create false/unnecessary state in the ENRP servers when the replay is for registration/de-registration or update.

### 2.8.2. Effect

The result is that false/extra state is maintained by ENRP servers. This would most likely be used as a denial-of-service attack if the replay is used to de-register all PEs.

### 2.8.3. Requirement

The protocol MUST prevent replay attacks. The replay attack prevention mechanism in TLS meets this requirement.

## 2.9. Re-Establishing PU-PE Security during Failover

### 2.9.1. Threat

The PU fails over from PE A to PE B. In the case that the PU had a trusted relationship with PE A, the PU will likely not have the same relationship established with PE B.

### 2.9.2. Effect

If there was a trust relationship involving security context between PU and PE A, the equivalent trust relationship will not exist between PU and PE B. This will violate security policy. For example, if the security context with A involves encryption and the security context with B does not, then an attacker could take advantage of the change in security.

### 2.9.3. Requirement

The application SHOULD be notified when failover occurs so the application can take appropriate action to establish a trusted relationship with PE B. ENRP has a mechanism to perform this function.

## 2.10. Integrity

### 2.10.1. Threat

The following are all instances of the same class of threats, and all have similar effects.

- a. ENRP response to pool handle resolution is corrupted during transmission.
- b. ENRP peer messages are corrupted during transmission.
- c. PE sends an update for values, and that information is corrupted during transmission.

### 2.10.2. Effect

The result is that ASAP receives corrupt information for pool handle resolution, which the PU believes to be accurate. This corrupt information could be an IP address that does not resolve to a PE so the PU would not be able to contact the server.

### 2.10.3. Requirement

An integrity mechanism MUST be present. Corruption of data that is passed to the PU means that the PU can't rely on it. The consequence of corrupted information is that the IP addresses passed to the PU might be wrong, in which case, it will not be able to reach the PE. The interfaces that MUST implement integrity are PE to ENRP server and ENRP to ENRP server. The integrity mechanism in TLS is used for this.

## 2.11. Data Confidentiality

### 2.11.1. Threat

An eavesdropper capable of snooping on fields within messages in transit may be able to gather information, such as topology/location/IP addresses, etc., which may not be desirable to divulge.

### 2.11.2. Effect

Information that an administrator does not wish to divulge is divulged. The attacker gains valuable information that can be used for financial gain or attacks on hosts.

### 2.11.3. Requirement

A provision for data confidentiality service SHOULD be available. TLS provides data confidentiality in support of this mechanism.

## 2.12. ENRP Server Discovery

### 2.12.1. Threats

- a. Thwarting successful discovery: When a PE wishes to register with an ENRP server, it needs to discover an ENRP server. An attacker could thwart the successful discovery of ENRP server(s), thereby inducing the PE to believe that no ENRP server is available. For instance, the attacker could reduce the returned set of ENRP servers to null or a small set of inactive ENRP servers. The attacker performs a MITM attack to do this.
- b. A similar thwarting scenario also applies when an ENRP server or ASAP on behalf of a PU needs to discover ENRP servers.
- c. Spoofing successful discovery: An attacker could spoof the discovery by claiming to be a legitimate ENRP server. When a PE wishes to register, it finds the spoofed ENRP server. An attacker can only make such a claim if no security mechanisms are used.
- d. A similar spoofing scenario also applies when an ENRP server or ASAP on behalf of a PU needs to discover ENRP servers.

### 2.12.2. Effects (Letters Correlate with Threats above)

- a. A PE that could have been in an application server pool does not become part of a pool. The PE does not complete discovery operation. This is a DoS attack.
- b. An ENRP server that could have been in an ENRP server pool does not become part of a pool. A PU is unable to utilize services of ENRP servers.
- c. This malicious ENRP would either misrepresent, ignore, or otherwise hide or distort information about the PE to subvert RSerPool operation.
- d. Same as above.

### 2.12.3. Requirement

A provision for authentication **MUST** be present and a provision for data confidentiality service **SHOULD** be present. TLS has a mechanism for confidentiality.

## 2.13. Flood of Endpoint-Unreachable Messages from the PU to the ENRP Server

### 2.13.1. Threat

Endpoint-unreachable messages are sent by ASAP to the ENRP server when it is unable to contact a PE. There is the potential that a PU could flood the ENRP server intentionally or unintentionally with these messages. The non-malicious case would require an incorrect implementation. The malicious case would be caused by writing code to flood the ENRP server with endpoint unreachable messages.

### 2.13.2. Effect

The result is a DoS attack on the ENRP server. The ENRP server would not be able to service other PUs effectively and would not be able to take registrations from PEs in a timely manner. Further, it would not be able to communicate with other ENRP servers in the pool to update the database in a timely fashion.

### 2.13.3. Requirement

The number of endpoint unreachable messages sent to the ENRP server from the PU **SHOULD** be limited. This mechanism is described in the ASAP [RFC5352] protocol document.

## 2.14. Flood of Endpoint Keep-Alive Messages from the ENRP Server to a PE

### 2.14.1. Threat

Endpoint Keep-Alive messages would be sent from the ENRP server to the PEs during the process of changing the Home ENRP server for this PE.

### 2.14.2. Effect

If the ENRP server maliciously sent a flood of endpoint Keep-Alive messages to the PE, the PE would not be able to service clients. The result is a DoS attack on the PE.

### 2.14.3. Requirement

ENRP MUST limit the frequency of Keep-Alive messages to a given PE to prevent overwhelming the PE. This mechanism is described in the ENRP [RFC5353] protocol document.

## 2.15. Security of the ENRP Database

### 2.15.1. Threat

Another consideration involves the security characteristics of the ENRP database. Suppose that some of the PEs register with an ENRP server using security and some do not. In this case, when a client requests handlespace resolution information from ENRP, it would have to be informed which entries are "secure" and which are not.

### 2.15.2. Effect

This would not only complicate the protocol, but actually bring into question the security and integrity of such a database. What can be asserted about the security of such a database is a very thorny question.

### 2.15.3. Requirement

The requirement is that either the entire ENRP server database MUST be secure; that is, it has registrations exclusively from PEs that have used security mechanisms, or the entire database MUST be insecure; that is, registrations are from PEs that have used no security mechanisms. ENRP servers that support security MUST reject any PE server registration that does not use the security mechanisms. Likewise, ENRP servers that support security MUST NOT accept updates from other ENRP servers that do not use security mechanisms. TLS is used as the security mechanism so any information not sent using TLS to a secure ENRP server MUST be rejected.

## 2.16. Cookie Mechanism Security

The application layer is out of scope for RSerPool. However, some questions have been raised about the security of the cookie mechanism, which will be addressed.

Cookies are passed via the ASAP control channel. If TCP is selected as the transport, then the data and control channel MUST be multiplexed. Therefore, the cases:

- a. control channel is secured; data channel is not

b. data channel is secured; control channel is not

are not possible, as the multiplexing onto one TCP port results in security for both data and control channels or neither.

The multiplexing requirement results in the following cases:

1. the multiplexed control channel-data channel is secure; \*or\*
2. the multiplexed control channel-data channel is not secured.

This applies to cookies in the sense that, if you choose to secure your control-data channel, then the cookies are secured.

A second issue is that the PE could choose to sign and/or encrypt the cookie. In this case, it must share keys and other information with other PEs. This application-level state sharing is out of scope of RSerPool.

#### 2.17. Potential Insider Attacks from Legitimate ENRP Servers

The previous text does not address all byzantine attacks that could arise from legitimate ENRP servers. True protection against misbehavior by authentic (but rogue) servers is beyond the capability of TLS security mechanisms. Authentication using TLS does not protect against byzantine attacks, as authenticated ENRP servers might have been maliciously hacked. Protections against insider attacks are generally specific to the attack, so more experimentation is needed. For example, the following discusses two insider attacks and potential mitigations.

One issue is that legitimate users may choose not to follow the proposed policies regarding the choice of servers (namely, members in the pool). If the "choose a member at random" policy is employed, then a pool user can always set its "random choices" so that it picks a particular pool member. This bypasses the "load sharing" idea behind the policy. Another issue is that a pool member (or server) may report a wrong policy to a user.

To mitigate the first attack, the protocol may require the pool user to "prove" to the pool member that the pool member was chosen "randomly", say by demonstrating that the random choice was the result of applying some hash function to a public nonce. Different methods may be appropriate for other member scheduling policies.

To mitigate the second attack, the protocol might require the PE to sign the policy sent to the ENRP server. During pool handle resolution, the signed policy needs to be sent from an ENRP server to an ASAP endpoint in a way that will allow the user to later hold the server accountable to the policy.

### 3. Security Considerations

This informational document characterizes potential security threats targeting the RSerPool architecture. The security mechanisms required to mitigate these threats are summarized for each architectural component. It will be noted which mechanisms are required and which are optional.

From the threats described in this document, the security services required for the RSerPool protocol suite are given in the following table.

| Threat       | Security mechanism in response  |
|--------------|---|
| Section 2.1  | ENRP server authenticates the PE.   |
| Section 2.2  | ENRP server authenticates the PE.   |
| Section 2.3  | ENRP server authenticates the PE.   |
| Section 2.4  | ENRP server authenticates the PE.   |
| Section 2.5  | ENRP servers mutually authenticate.   |
| Section 2.6  | PE authenticates the ENRP server.   |
| Section 2.7  | The PU authenticates the ENRP server. If the authentication fails, it looks for another ENRP server.  |
| Section 2.8  | Security protocol that has protection from replay attacks.  |
| Section 2.9  | Either notify the application when failover occurs so the application can take appropriate action to establish a trusted relationship with PE B *or* re-establish the security context transparently. |
| Section 2.10 | Security protocol that supports integrity protection.   |
| Section 2.12 | Security protocol that supports data confidentiality.   |
| Section 2.11 | The PU authenticates the ENRP server. If the authentication fails, it looks for another ENRP server.  |
| Section 2.13 | ASAP must control the number of endpoint unreachable messages transmitted from the PU to the ENRP server.   |
| Section 2.14 | ENRP server must control the number of Endpoint_KeepAlive messages to the PE.   |

The first four threats, combined with the sixth threat, result in a requirement for mutual authentication of the ENRP server and the PE.

To summarize, the first twelve threats require security mechanisms that support authentication, integrity, data confidentiality, and protection from replay attacks. For RSerPool, we need to authenticate the following:

- o PU -----> ENRP Server (PU authenticates the ENRP server)
- o PE <-----> ENRP Server (mutual authentication)
- o ENRP server <-----> ENRP Server (mutual authentication)



## Summary by component:

RSerPool client -- mandatory-to-implement authentication of the ENRP server is required for accurate pool handle resolution. This is to protect against threats from rogue ENRP servers. In addition, confidentiality, integrity, and preventing replay attack are also mandatory to implement to protect from eavesdropping and data corruption or false data transmission. Confidentiality is mandatory to implement and is used when privacy is required.

PE to ENRP communications -- mandatory-to-implement mutual authentication, integrity, and protection from replay attack is required for PE to ENRP communications. This is to protect the integrity of the ENRP handlespace database. Confidentiality is mandatory to implement and is used when privacy is required.

ENRP to ENRP communications -- mandatory-to-implement mutual authentication, integrity, and protection from replay attack is required for ENRP to ENRP communications. This is to protect the integrity of the ENRP handlespace database. Confidentiality is mandatory to implement and is used when privacy is required.

## 4. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.
- [RFC5352] Stewart, R., Xie, Q., Stillman, M., and M. Tuexen, "Aggregate Server Access Protocol (ASAP)", RFC 5352, September 2008.
- [RFC5353] Xie, Q., Stewart, R., Stillman, M., Tuexen, M., and A. Silverton, "Endpoint Handlespace Redundancy Protocol (ENRP)", RFC 5353, September 2008.
- [RFC5351] Lei, P., Ong, L., Tuexen, M., and T. Dreibholz, "An Overview of Reliable Server Pooling Protocols", RFC 5351, September 2008.

## Authors' Addresses

Maureen Stillman, Ed.  
Nokia  
1167 Peachtree Court  
Naperville, IL 60540  
USA

E-Mail: [maureen.stillman@nokia.com](mailto:maureen.stillman@nokia.com)

Ram Gopal  
Nokia Siemens Networks  
12278 Scripps Summit Drive  
San Diego, CA 92131  
USA

E-Mail: [ram.gopal@nsn.com](mailto:ram.gopal@nsn.com)

Erik Guttman  
Sun Microsystems  
Eichhoelzelstrasse 7  
74915 Waibstadt  
DE

E-Mail: [Erik.Guttman@sun.com](mailto:Erik.Guttman@sun.com)

Senthil Sengodan  
Nokia Siemens Networks  
6000 Connection Drive  
Irving, TX 75039  
USA

E-Mail: [Senthil.sengodan@nsn.com](mailto:Senthil.sengodan@nsn.com)

Matt Holdrege

E-Mail: [Holdrege@gmail.com](mailto:Holdrege@gmail.com)

## Full Copyright Statement

Copyright (C) The IETF Trust (2008).

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.

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, THE IETF TRUST 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.

## Intellectual Property

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 procedures with respect to rights in RFC 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).