Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 31 July 2026

Z. Luo, Ed.
H. Yan
CMCC

27 January 2026

Available Session Recovery Protocol
draft-cmcc-asrp-05

Abstract

This document describes an experimental protocol named the Available Session Recovery Protocol (ASRP). The protocol is designed to optimize high-availability network cluster architectures, providing a superior high-availability solution for clusters offering stateful network services such as load balancing and Network Address Translation (NAT [RFC4787]). ASRP defines the procedures for session backup and recovery, as well as the message formats used during these interactions, enabling efficient and streamlined session state management.

In contrast to traditional high-availability techniques that back up session state within the cluster itself, the core innovation of ASRP lies in its distributed backup of state information to the client or server side. This approach offers multiple advantages: it significantly enhances the cluster’s elastic scaling capabilities; supports rapid recovery from single-point or even multi-point failures; reduces resource redundancy by eliminating centralized backup nodes; and substantially simplifies the implementation complexity of the cluster.

The ASRP protocol provides exceptional elastic scalability for network clusters, facilitating the implementation and deployment of large-scale elastic network clusters.

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/.

Copyright Notice

Table of Contents

1. Introduction

Traditional high-availability network clusters based on a master-backup architecture rely on session state synchronization between the master and backup nodes. While functionally complete, this architecture faces challenges in the cloud era, such as insufficient flexibility for elastic scaling, resource redundancy, and high implementation complexity. To address these challenges, the industry has proposed the Elastic Stateful Cluster.

An Elastic Stateful Cluster is a high-availability network service cluster composed of multiple cooperative nodes. The number of nodes within the cluster can be elastically scaled, enabling it to provide stateful network services such as load balancing (SLB) and Network Address Translation (NAT). To achieve elastic scaling, conventional Elastic Stateful Clusters adopt a Fast/Slow Path design philosophy, separating session management from packet forwarding. This allows the fast path node layer to achieve good elastic scaling capabilities.

1.1. Conventional Elastic Stateful Cluster

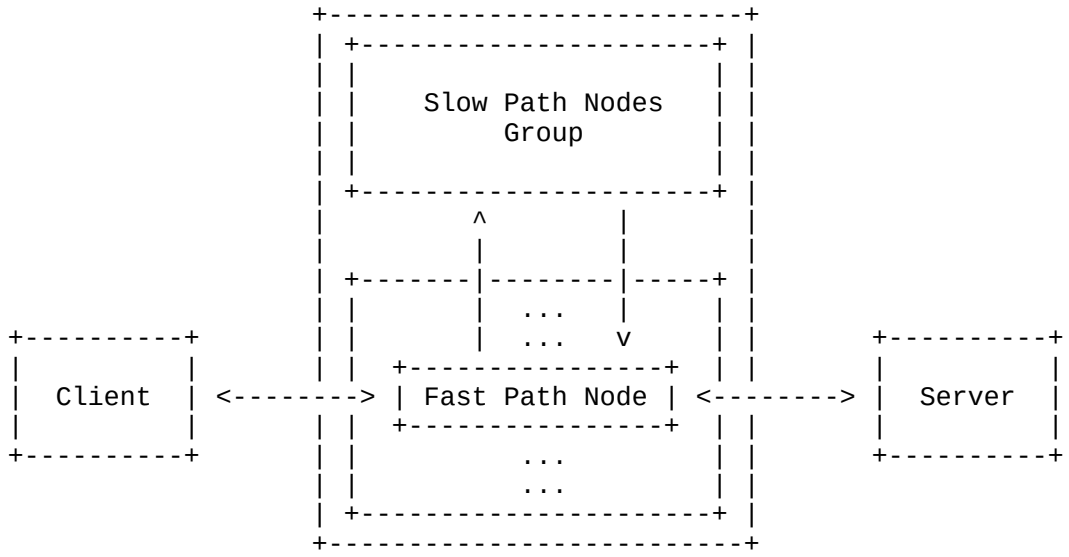

Figure 1: Fast/Slow Path Elastic Stateful Cluster

The slow path nodes are responsible for session creation and synchronization, while the fast path nodes are responsible for rapid packet forwarding. The drawback of this Elastic Stateful Cluster architecture is the weak elastic scaling capability of the slow path nodes. Implementing session synchronization among slow path nodes is complex. A typical implementation reference is the AWS Hyperplane NFV platform.

1.2. ASRP Elastic Stateful Cluster

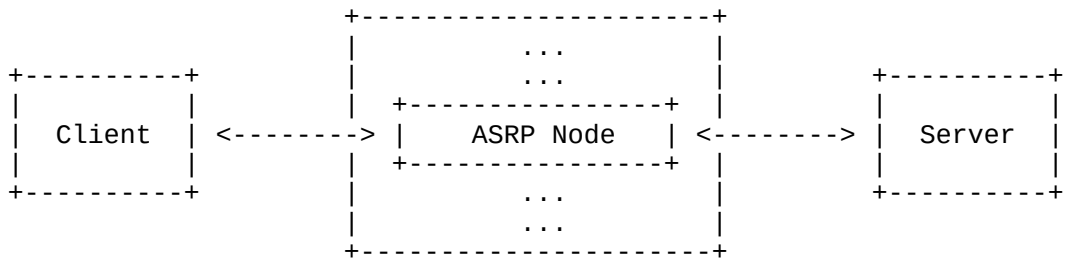

Figure 2: ASRP Elastic Stateful Cluster

The Available Session Recovery Protocol (ASRP) proposes an innovative high-availability solution aimed at building a more concise, efficiently elastic, and highly available architecture for stateful services. Its core idea is to innovatively distribute session state

information to the client or server. The lifecycle of the backup state is synchronized with the real session, eliminating the need for independent keepalive and timeout mechanisms. This design ensures the timeliness and availability of the backup information.

ASRP defines corresponding session backup and recovery mechanisms. The protocol allows protocol messages to be transmitted together with the original service data packets, thereby reducing control overhead for state synchronization. In an elastic stateful cluster built on ASRP, network nodes possess atomic and mutually independent properties. There is no need for communication between nodes, nor is session synchronization required within the cluster. This fundamental design provides theoretically unlimited scaling capability and supports rapid recovery from single-point or even multi-point failures.

2. Terminology

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 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Protocol Overview

3.1. Two Operational Modes

For the ASRP protocol to function correctly, two prerequisites must be met. First, all network nodes within the cluster MUST run service software supporting the ASRP protocol. Second, the server or client responsible for backing up sessions MUST deploy a kernel module or an eBPF module that supports ASRP. Depending on whether this module is deployed on the server or the client, the protocol operates in one of two corresponding modes: Passive (PSV) Mode and Active (ACT) Mode.

3.1.1. PSV Mode

In PSV mode, the network node is typically located within the same trusted network domain as the server (e.g., inside a data center). Its typical service is load balancing.

3.1.2. ACT Mode

In ACT mode, the network node is typically located within the same trusted network domain as the client (e.g., an enterprise intranet). Its typical service is Source Network Address Translation (SNAT).

3.2. Two Routing Behaviors

3.2.1. Symmetric Routing

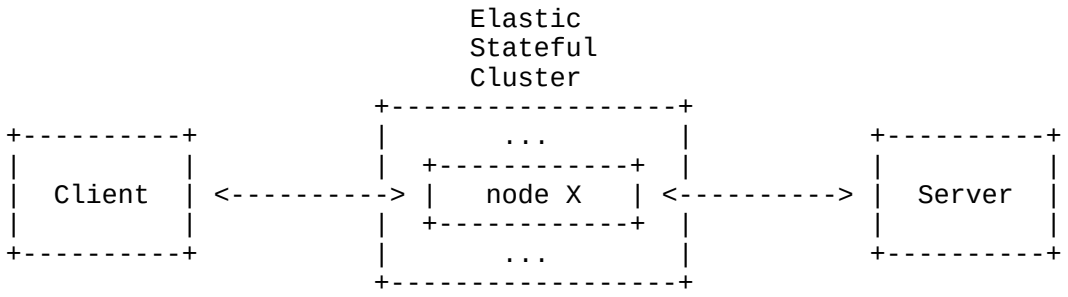

Figure 3: Symmetric Routing

Symmetric routing refers to the path mode where bidirectional traffic of the same session between a client and a server is always routed to the same node within the cluster.

3.2.2. Asymmetric Routing

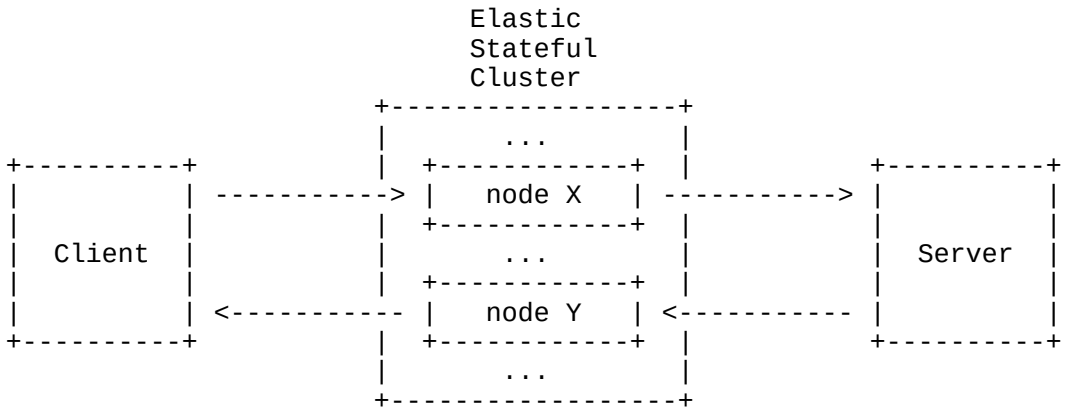

Figure 4: Asymmetric Routing

Asymmetric routing refers to the scenario where bidirectional traffic of the same session may be routed (e.g., by mechanisms such as ECMP [RFC2991], [RFC2992]) to different nodes within a cluster. In cloud networking environments, asymmetric routing is a common phenomenon, which imposes higher demands on the implementation of elastic stateful clusters.

3.3. Protocol Message

ASRP achieves distributed backup and recovery of session state information by exchanging specific protocol messages among the client, server, and network nodes (such as load balancers or NAT devices). In a load-balancing scenario, session state is distributed and backed up to individual servers; in a Source Network Address Translation (SNAT) scenario, session state is distributed and backed up to individual clients.

ASRP defines five protocol messages: New Session message (NS), Query Session message (QS), Recover Session message (RS), Hello Session message (HS), and Push Session Message (PS). NS, QS, and RS messages are encapsulated within UDP (not UDP-lite, [RFC0768], [RFC3828]) datagrams for transmission. A specific destination port, referred to as ASRP-PORT (currently a configurable experimental port, e.g., 51200, is used), identifies that the UDP payload contains an ASRP message. HS/PS messages adopt the same outer encapsulation as the forwarded packet (to ensure HS/PS packets are routed to the correct network node), employing an ASPR Signature to identify an ASRP message.

A packet carrying an ASRP message is termed an ASRP packet (NS/QS/RS/HS/PS packet). An ASRP packet can simultaneously carry both the ASRP message and the forwarded packet. If it carries only the ASRP message, it is referred to as a pure ASRP packet (pure NS/QS/RS/HS/PS packet).

3.3.1. NS Message

Generated by the network node, it is used to send session state information to a designated client (in ACT mode) or server (in PSV mode) for backup when creating a new session.

3.3.2. QS Message

Generated by the network node, it is used to query the client or server for backup session state information when a received packet cannot match any local session and a session cannot be directly created. For TCP SYN packets, if no local session matches, a session can be created directly without querying the state.

3.3.3. RS Message

Generated by the client or server holding the backup as a response to a NS/QS message, it contains the state information required to recover the session. The network node parses the RS message and reconstructs or marks the local session, thereby achieving failure recovery.

3.3.4. HS Message

Generated by the client, it is used in ACT mode to announce to the network node its capability to support the ASRP protocol and to trigger the network node to return an NS message to complete session backup.

3.3.5. PS Message

Generated by the server, it is used in PSV mode to push session state information to the network node. In the case of asymmetric routing, the network node utilizes the PS message to create/update sessions for fast packet forwarding.

3.4. Session Creation/Recovery Scenarios

This section elaborates on, through a series of typical scenarios, how the ASRP protocol achieves session backup and recovery via message interaction in the event of network node failures under different operational modes. Each scenario details the involved protocol message flows and the processing steps of each entity.

3.4.1. PSV-Scenario-1

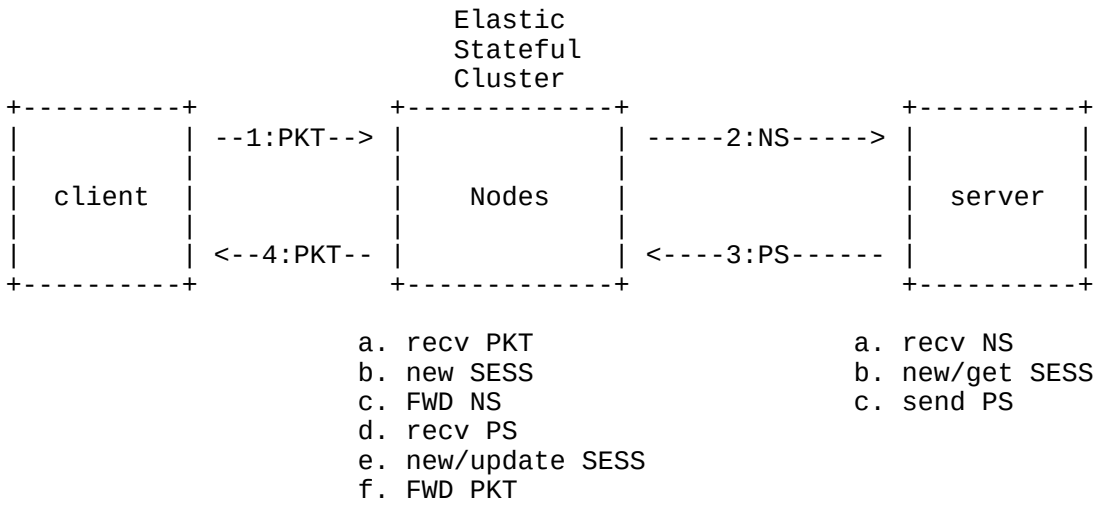

Figure 5: Direct Session Creation in PSV Mdoe

This scenario describes the direct session creation flow in PSV mode. The most common example is the SYN packet during TCP connection establishment, which represents the client initiating a new connection.

The processing flow is as follows:

1. Session Creation: Upon receiving a packet from a client (e.g., a TCP SYN), if no local session is found, the network node directly creates a new session. Subsequently, the network node sends an NS message to the selected server. If the NS message and the forwarded packet are transmitted separately, the NS message is sent first.

2. Server Response: Upon receiving the NS message, the server backs up the session state information contained in the NS message locally and associates it with its local session. In the case of asymmetric routing, when the server sends its first response packet, it generates an PS message and sends it to the network node.

3. Session Recovery: In the case of asymmetric routing, the network node, upon receiving the PS message, recovers the local session and subsequently forwards packets according to that session.

In the scenario above, provided that no IP fragmentation occurs, the NS/PS messages and the forwarded packets SHOULD be transmitted together to improve transmission efficiency. For example, for a TCP

session, NS/PS messages are generally transmitted together with the SYN/SYN-ACK ([RFC9293]) packets. The backed-up session state information is released when the local session closes, requiring no additional close messages

3.4.2. PSV-Scenario-2

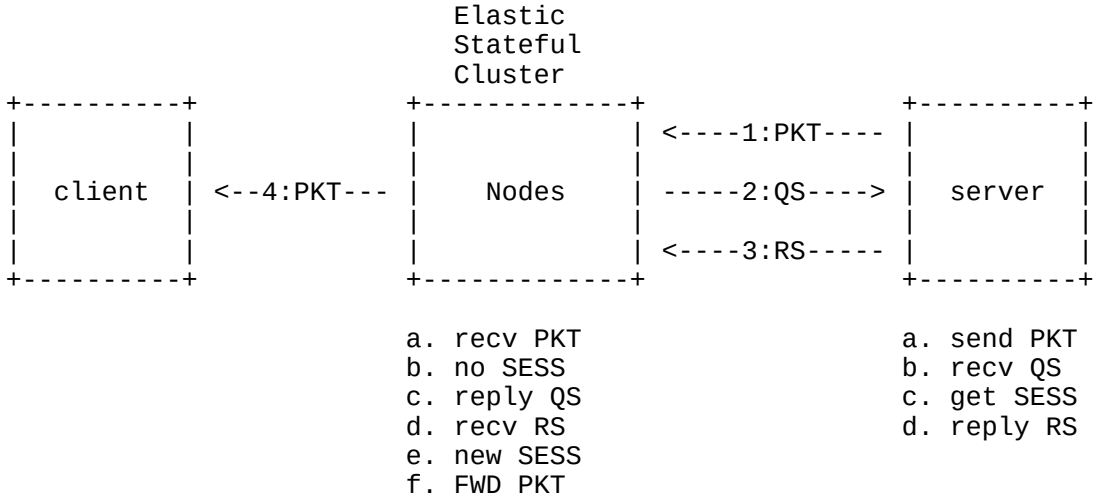

Figure 6: Session Recovery for Server in PSV Mode

This scenario describes the session recovery flow triggered by a server packet when the network node has lost the session in PSV mode.

The processing flow is as follows:

1. Session Query: Upon receiving a packet from the server, the network node searches its local session table. If no corresponding session is found, the network node generates a QS message and sends it back to the server.

2. Server-Assisted Reply: After receiving the QS message, the server, based on the content of the QS message, looks up the locally stored backup session state information and then generates an RS message, sending it back to the network node.

3. Session Recovery: After receiving the RS message, the network node creates a new local session and subsequently forwards packets according to that session.

In the scenario above, provided that no IP fragmentation would occur, the forwarded packet may either be buffered or transmitted together with the QS message; otherwise, the network node should buffer the forwarded packet. Once the session is recovered, any buffered packet MUST be processed immediately and forwarded in accordance with the session.

3.4.3. PSV-Scenario-3

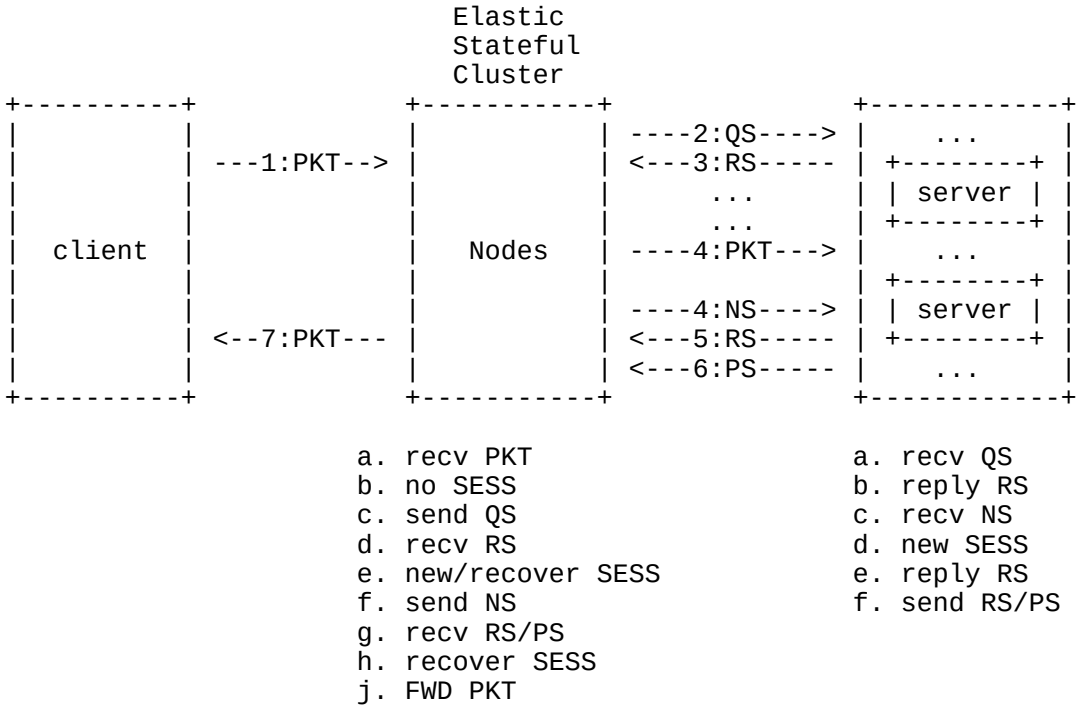

Figure 7: Session Creation/Recovery for Client in PSV Mode

This scenario describes the situation in PSV mode where, upon receiving a packet from a client, the network node cannot match it to a local session and cannot directly create a new session either. The network node MUST first determine whether this packet belongs to an existing session to decide how to handle it. The network node uses the ASRP protocol to query servers that may hold the backup session state information. ASRP relies on the cluster employing a deterministic server selection algorithm (such as a consistent hashing algorithm or a consistent hashing algorithm with history) to identify the target servers for querying.

A consistent hashing algorithm with history maintains a list of servers that have been used historically within a hash bucket. This list also serves as the target candidate server list for the network node's queries. ASRP recommends setting a maximum query count to avoid performance issues. Simultaneously, ASRP suggests setting a timeout for historical servers in the hash bucket to reduce the length of the server list by deleting timed-out historical records.

The processing flow is as follows:

1. Query Local Session: Upon receiving a forwarded packet from a client, the network node searches its local session table. If no local session is found, it calculates candidate servers (which may be multiple) using a deterministic server selection algorithm.

2. Query Backup Session: The network node sends QS messages to each candidate server to query for the backup session. The servers return the query results via RS messages.

3. Process Query Results: If a session is found, the network node recovers the local session based on the RS message and then forwards the forwarded packet. If no session is found, it proceeds to the new session creation flow by sending an NS message to the server selected by the algorithm.

4. Server Creates New Session: After receiving the NS message, the server backs up the session state information locally. In an asymmetric routing environment, it MUST immediately reply with a pure RS packet as an acknowledgment.

5. Session Recovery: When the server sends its first response packet to the client, it generates an PS message and sends it to the network node. Upon receiving the PS message, the network node first recovers the local session based on the message and then forwards packets according to the session.

In the scenario above, provided that no IP fragmentation would occur, the forwarded packet may either be buffered or transmitted together with the QS message; otherwise, the network node SHOULD buffer the forwarded packet. Once the session is created or recovered, any buffered packet MUST be processed immediately and forwarded in accordance with the session.

3.4.4. ACT-Scenario-1

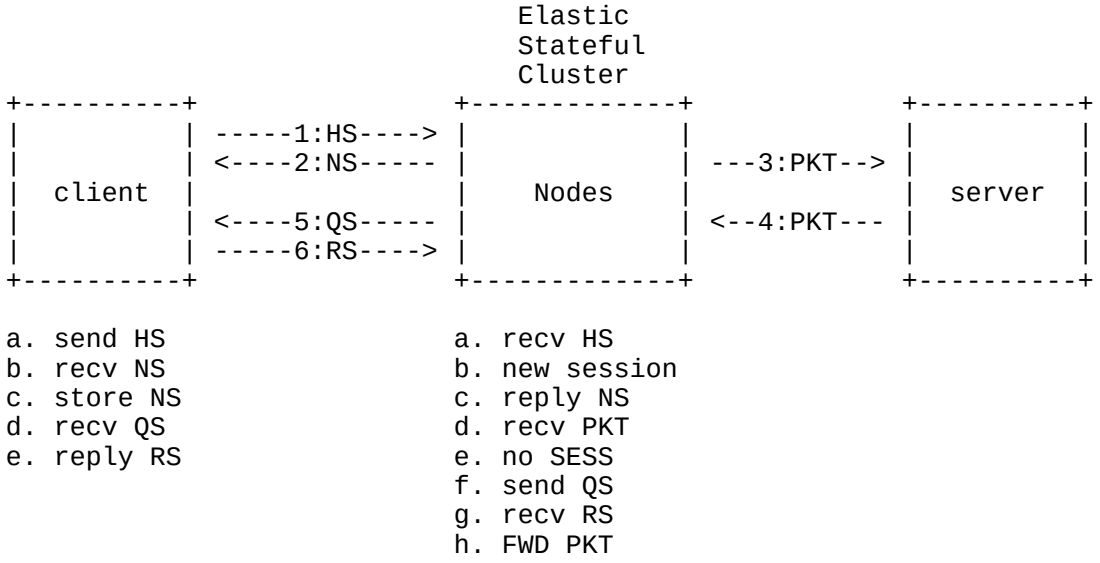

Figure 8: Session Creation/Recovery in ACT Mode

This scenario describes the session creation process by the network node and the server-packet-triggered session recovery flow in ACT mode. During the session recovery phase, the network node MUST be able to deterministically locate the client that holds the backup for that session. The use of a static, configurable mapping strategy is recommended. If such a mapping cannot be established, ASRP cannot function in this scenario. For SNAT services, ports can typically be used to map clients, with different clients using different, configurable port ranges.

The processing flow is as follows:

1. Session Creation: When a client initiates the first packet, it generates an HS message and sends it to the network node. Upon receiving the HS message, the network node follows the normal procedure to create a new session, returns a pure NS packet to the client, and forwards the forwarded packet according to the session.

2. Processing Server Response Packets: If a matching session is found, the packet is forwarded according to that session. If no matching session is found, the network node uses the mapping relationship to locate the corresponding client and sends a QS message to it.

3. Client-Assisted Recovery: After receiving the QS message, the client queries its locally stored backup session state information and replies with an RS message to the network node.

4. Session Recovery: After receiving the RS message, the network node recovers the session state locally and subsequently forwards packets according to the session.

After sending an HS message, the client waits for an NS message. If an NS message is not received, a minimum time interval (suggested on the order of milliseconds) is set. Subsequent packets sent by the client will trigger new HS messages to remind the network node to return an NS message. Upon receiving an HS message, if the network node does not find a matching local session, it creates a session, generates an NS message, and sends it to the client. If the network node subsequently receives further HS messages that do match a local session, it will also immediately send an NS message to the client.

In the scenario above, provided that no IP fragmentation would occur, the forwarded packet may either be buffered or transmitted together with the QS message; otherwise, the network node SHOULD buffer the forwarded packet. Once the session is recovered, any buffered packet MUST be processed immediately and forwarded in accordance with the session.

3.4.5. ACT-Scenario-2

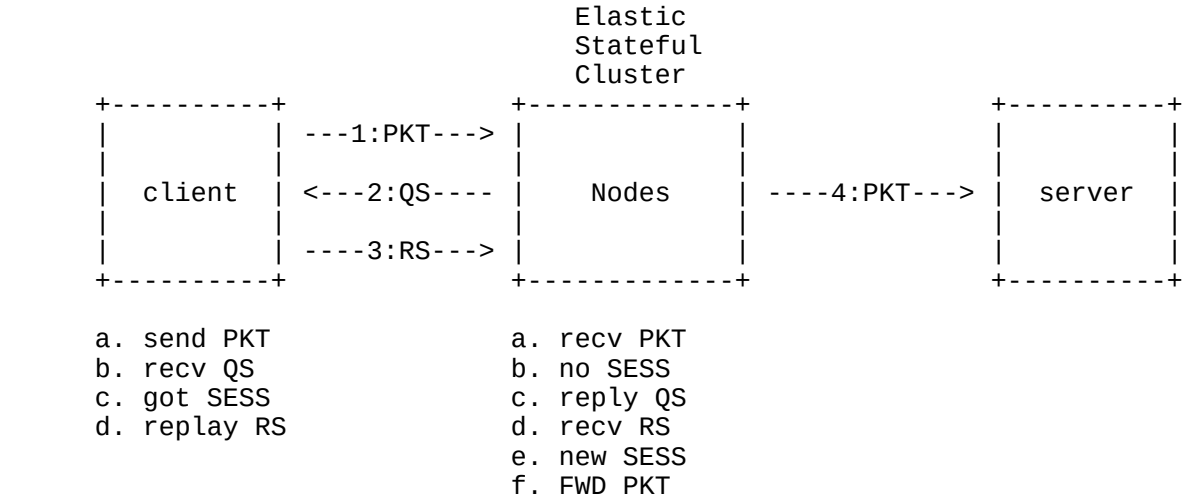

Figure 9: Session Recovery for Client in ACT Mode

This scenario describes the client-packet-triggered session recovery flow in ACT mode.

The processing flow is as follows:

1. Session Query: Upon receiving a packet from a client, if the network node finds no local session and the packet does not contain an HS message, it sends a QS message to the client.

2. Client-Assisted Recovery: After receiving the QS message, the client queries its locally stored backup session state information and replies with an RS message to the network node.

3. Session Recovery: After receiving the RS message, the network node recovers the session locally and subsequently forwards the packet according to the session.

In the scenario above, provided that no IP fragmentation would occur, the forwarded packet may either be buffered or transmitted together with the QS message; otherwise, the network node SHOULD buffer the forwarded packet. Once the session is recovered, any buffered packet MUST be processed immediately and forwarded in accordance with the session.

4. Protocol Details

4.1. Message Format

All ASRP protocol messages are encoded using the TLV (Type-Length-Value) structure. All numeric fields use network byte order (big-endian).

The fields that can be used in ASRP messages are as follows:

1. Sub and Type: 1 byte. Sub (high 4 bits) indicates the internal data type of the message; Type (low 4 bits) indicates the message type.

2. Length: 1 byte, indicating the total length of the entire ASRP message.

3. Flags: 1 byte. ASRP_F_ACT (0x1) is ACT mode flag; ASRP_F_MSG (0x2) is pure message flag.

4. Protocol: 1 byte, identifying the session protocol, such as TCP, UDP, etc.

5. Session-Tuple: Contains source address, destination address, source port, destination port. The IP address type is IPv4 or IPv6 ([RFC0791] [RFC8200]).

6. Session-Data: Variable-length field, carrying the private state information of the network node. The specific content is determined by the implementation and can generally be empty.

If the ASRP_F_ACT flag is set, it indicates the current mode is ACT mode; otherwise, the current mode is PSV mode.

If the ASRP_F_MSG flag is set, it indicates the message is transmitted independently; otherwise, it indicates this message is transmitted together with the forwarded packet.

IPv4-Session-Tuple Format:

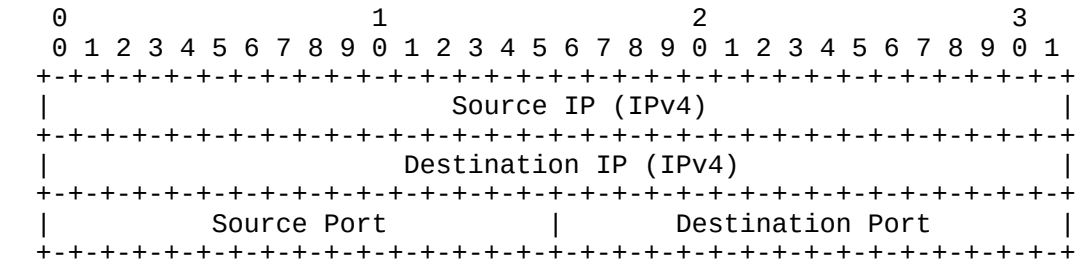

Figure 10: IPv4 Session Tuple Format

IPv6-Session-Tuple Format: The structure of IPv6-Session-Tuple is the same as IPv4-Session-Tuple, with the main difference being the IP address length/type in the Session-Tuple field.

4.1.1. NS Message Format

The NS message is used by the network node to back up session state information to the client or server. The NS message contains two Session-Tuples.

Type Assignments (Least significant nibble):
NS: 0x0

Sub Assignments (Most significant nibble):
ST44: 0x0, IPv4-Session-Tuple + IPv4-Session-Tuple
ST66: 0x1, IPv6-Session-Tuple + IPv6-Session-Tuple
ST46: 0x2, IPv4-Session-Tuple + IPv6-Session-Tuple
ST64: 0x3, IPv6-Session-Tuple + IPv4-Session-Tuple

NS(Sub-ST44/ST66/ST46/ST64) Message Format:

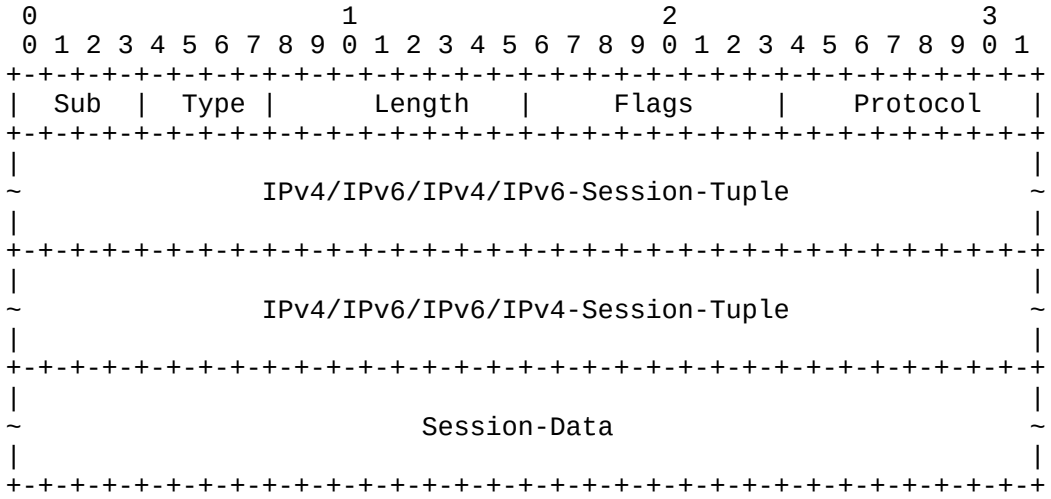

Figure 11: ASRP NS Message Format

The NS message contains two Session-Tuples, representing the connection between the network node and the client, and the connection between the network node and the server, respectively.

4.1.2. QS Message Format

The QS message is used by the network node to query backup session state information.

Type Assignments (Least significant nibble):
QS: 0x1

Sub Assignments (Most significant nibble):
ST4: 0x0, IPv4-Session-Tuple
ST6: 0x1, IPv6-Session-Tuple

QS(Sub-ST4/ST6) Message Format:

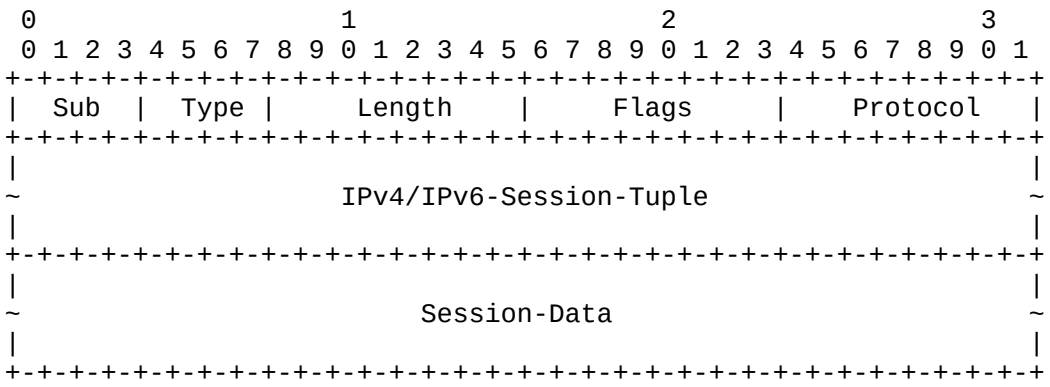

Figure 12: ASRP QS Message Format

4.1.3. RS Message Format

The RS message is used to recover a network node’s session.

Type Assignments (Least significant nibble):
RS: 0x2

Sub Assignments (Most significant nibble):
ST44: 0x0, IPv4-Session-Tuple + IPv4-Session-Tuple
ST66: 0x1, IPv6-Session-Tuple + IPv6-Session-Tuple
ST46: 0x2, IPv4-Session-Tuple + IPv6-Session-Tuple
ST64: 0x3, IPv6-Session-Tuple + IPv4-Session-Tuple
ST4: 0x4, IPv4-Session-Tuple
ST6: 0x5, IPv6-Session-Tuple

RS(Sub-ST44/ST66/ST46/ST64) Message Format: The structure of messages is the same as NS.

RS(Sub-ST4/ST6) Message Format: The structure of messages is the same as QS.

If the Sub field of an RS message is ST44, ST66, ST46, or ST64, it indicates the RS message carries session recovery information. If the Sub field is ST4 or ST6, this RS message is a response indicating a failed query for the corresponding QS message.

4.1.4. HS Message Format

The HS message is generated by the client to announce to the network node that it requires an NS message to back up session state information.

Type Assignments (Least significant nibble):
HS: 0x3

Sub Assignments (Most significant nibble):
NST: 0x0, No-Session-Tuple

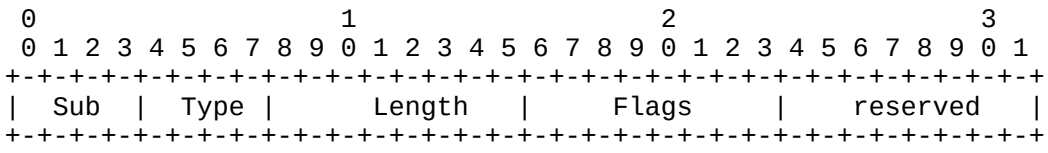

Figure 13: ASRP HS Message Format

4.1.5. PS Message Format

Type Assignments (Least significant nibble):
PS: 0x4

The structure of the PS message is the same as that of the NS message.

4.2. ASRP packet Format

4.2.1. NS/QS/RS packet

The format of the NS/QS/RS packet is as follows:

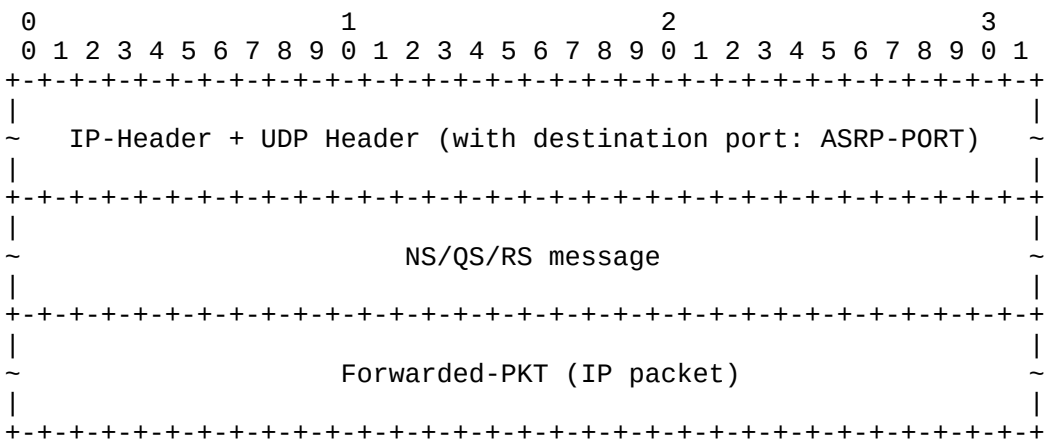

Figure 14: NS/QS/RS packet

If the ASRP_F_MSG flag is set in the Flags field of an NS/QS/RS message, it indicates that this is a pure NS/QS/RS packet (in which case the Forwarded-PKT section has a length of zero); otherwise, it indicates that the NS/QS/RS message is transmitted together with the forwarded packet.

4.2.2. HS/PS packet

HS/PS packets adopt the same protocol header as the forwarded packet (the packet sent from the client or server to the network node for forwarding)., which necessitates the use of an ASRP Signature to identify the message. Similar to the Proxy Protocol, a 12-byte ASRP Signature is placed at the beginning of the data: 0x0D 0x0A 0x0A 0x0D 0x00 0x0D 0x0A 0x41 0x53 0x52 0x50 0x0A. This signature contains a CRLF pattern, a null byte, and the specific ASCII sequence "ASRP". The probability of this sequence occurring in normal data streams is less than 2^{-96}, making it easy to debug and identify: during packet capture analysis, the clear "ASRP" identifier is visible.

The format of the HS/PS packet is as follows:

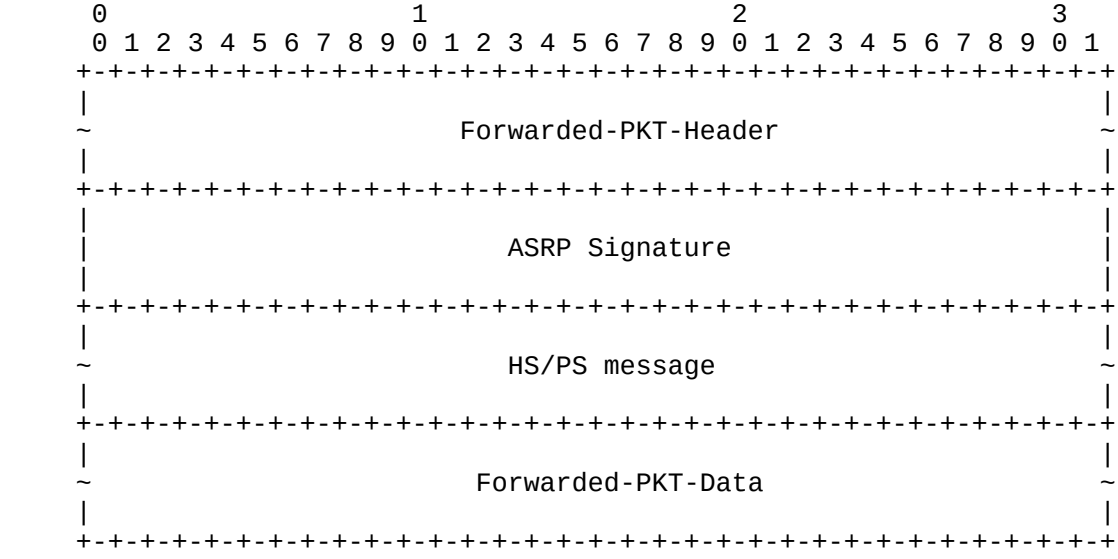

Figure 15: HS/PS packet

If the ASRP_F_MSG flag is set in the Flags field of an HS/PS message, it indicates that this is a pure HS/PS packet (in which case the Forwarded-PKT-Data section has a length of zero); otherwise, it indicates that HS/PS messages are embedded within the forwarded packet and transmitted together with it.

4.3. Message Processing

NS/QS/RS packets can be identified by their UDP destination port, while HS/PS packets are identified by the ASRP Signature. Once an ASRP packet is identified, the ASRP message within the packet can then be parsed and processed.

If transmission can be performed without causing IP fragmentation, all ASRP messages may be transmitted together with the forwarded packet. The specific encapsulation method is defined in the ASRP Packet Format. In subsequent message processing descriptions, this point will not be repeatedly emphasized.

4.3.1. NS Message Processing

The NS message is generated by a network node when creating a new session and is used to back up the session to the client or server.

The source IP of the NS packet is set to the network node’s local IP (which can be obtained from configuration), and the destination IP is set to the client’s or server’s IP (which can be obtained from the forwarded packet). The source port is randomly generated, and the destination port is set to ASRP-PORT.

When a client or server receives an NS packet, it extracts the NS message and backs up the session state information, extracts the forwarded packet (if present), and hands it over to the system.

In PSV mode, if an NS message is lost, for TCP connections, the retransmission of the SYN packet will trigger the retransmission of the NS message. For other types of connections, subsequent packets will continue to generate NS messages until an RS message is received.

In ACT mode, if an NS message is lost, subsequent packets sent by the client will generate HS messages, prompting the network node to retransmit the NS message in response to these subsequent HS messages.

NS messages may be generated in both PSV and ACT modes. The handling procedures are described in Figure 5, Figure 7, and Figure 8.

4.3.2. QS Message Processing

The QS message is generated by the network node to query backup session state information.

The source IP of the QS packet is set to the network node's local IP (obtainable from configuration), and the destination IP is set to the client's or server's IP (obtainable from the forwarded packet as described in Figure 6 and Figure 9, or derived via algorithmic mapping to the client or server as described in Figure 7 and Figure 8). The source port is randomly generated, and the destination port is set to ASRP-PORT.

When a client or server receives a QS packet, it extracts the QS message, queries the backup session state information, and returns an RS message; it extracts the forwarded packet (if present), processes it first according to the backup session state information, and then hands it over to the system.

If a QS message is lost, subsequent packets will trigger the generation of new QS packets, continuing the attempt to recover the session.

QS messages may be generated in both PSV and ACT modes. The handling procedures are described in Figure 6, Figure 7, and Figure 9.

4.3.3. RS Message Processing

The RS message is generated by the client or server in response to an NS or QS message. It is processed by the network node to recover a session.

The RS packet reuses the protocol header of the NS/QS packet, with the source and destination IP addresses and UDP ports swapped.

When a network node receives an RS packet, it extracts the RS message and recovers the session (upon successful QS query); it extracts the forwarded packet (if present) and forwards it according to the session.

If an RS message is lost, subsequent NS or QS messages will continue the attempt to recover the session, thereby triggering retransmission of the RS message.

RS messages may be generated in both PSV and ACT modes. The handling procedures are described in Figure 5, Figure 6, Figure 7, Figure 8, and Figure 9.

4.3.4. HS Message Processing

The HS message is sent by the client during the initial connection establishment phase to announce to the network node that it requires an NS message to back up session state information.

The source IP, destination IP, and source port of the HS packet are copied from the packet sent by the client.

When a network node receives an HS message, it extracts the HS message, creates a session, forwards packets according to the session, and returns a pure NS packet to the client.

HS messages are only generated in ACT mode. The handling procedure is described in Figure 8.

4.3.5. PS Message Processing

The PS message is generated by the server after receiving an NS message. When the server sends its first response packet to the client, it uses the PS message to push session state information to the network node. This is used to recover the network node’s session in the case of asymmetric routing.

The source IP, destination IP, and source port of the PS packet are copied from the packet sent by the server.

When a network node receives a PS message, it extracts the PS message and recovers the session; it extracts the forwarded packet (if present) and forwards it according to the session.

PS messages are only generated in PSV mode. The handling procedure is described in Figure 5 Figure 7.

5. Security Considerations

5.1. Message Forgery Attacks

The security design of the ASRP protocol is based on its typical deployment model.

Deployment Boundaries and Access Control: ASRP recommends deploying network nodes and the clients or servers that back up sessions within the same trusted internal network domain. In this model, all ASRP protocol packets communicate within an internal address space. By implementing appropriate network segmentation (e.g., using firewall policies or security groups) and strictly checking the source addresses of packets, forged ASRP packets originating from untrusted external networks can be effectively prevented from reaching the target nodes.

Session Legitimacy Verification: When processing ASRP packets that may establish new sessions (e.g., HS or RS packets), network nodes SHOULD perform basic validation according to the specific policies of

the upper-layer application or service. For instance, in a load-balancing scenario, a node SHOULD verify whether the session points to a known and healthy server. In a NAT scenario, it SHOULD verify whether the address translation complies with predefined rules. This prevents the establishment of illegal sessions at the application layer.

Internal Threat Assessment: Even if an attacker is located within the trusted network and can forge ASRP packets, the scope of impact is inherently limited. The attacker can only forge sessions where they themselves are the endpoint (e.g., masquerading as a client to request recovery of a non-existent connection). Such forged sessions are indistinguishable in form from sessions established through normal access. They do not directly jeopardize the security of other users or nodes, nor can they elevate the attacker’s privileges or grant access to unauthorized resources.

5.2. QS/RS Flood Attacks

When a network node loses a session, it may generate a large volume of QS packets. If maliciously exploited or due to a malfunction, this could lead to a flood attack [RFC4987]. To mitigate such risks, implementers SHOULD consider the following protective measures:

Rate Limiting and Traffic Shaping: Each network node SHOULD implement monitoring and limiting of the rate at which QS packets are sent. A reasonable threshold (e.g., the number of QS packets allowed per second) SHOULD be set. When the rate exceeds this threshold, the node SHOULD adopt a packet drop policy, for example, discarding newly arriving forwarded packets that trigger queries. The parameters for rate limiting SHOULD be configurable to adapt to deployment environments of different scales.

6. IANA Considerations

This document defines an application-layer protocol (ASRP). The protocol message types and internal identifiers are defined by this specification itself and constitute internal implementation details of the protocol. Therefore, there is no need to request registration of a separate protocol number or code point from IANA. However, for the implementation of this protocol, a UDP destination port requires allocation:

6.1. UDP Destination Port

NS/QS/RS messages are encapsulated within UDP datagrams for transmission. A fixed UDP destination port number is required so that the receiving end can identify and process such encapsulated packets.

Service Name: asrp

Port Number: 51200 (proposed value for current experimentation)

Transport Protocol: udp

Description: Used for receiving UDP-encapsulated ASRP protocol messages.

For experimental implementations and interoperability testing prior to IANA assignment, UDP port 51200 MAY be used as a temporary default. This port falls within the dynamic/private port range (49152-65535) reserved for local or temporary use and documentation examples [RFC6335].

IANA is requested to assign a permanent port number in the "User Ports" range (1024-49151) for the "asrp" service in the "Service Name and Transport Protocol Port Number Registry", with a reference to this document.

7. References

7.1. Normative References

[RFC0768] Postel, J., "User Datagram Protocol", STD 6, RFC 768, DOI 10.17487/RFC0768, August 1980, <https://www.rfc-editor.org/info/rfc768>.

[RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, DOI 10.17487/RFC0791, September 1981, <https://www.rfc-editor.org/info/rfc791>.

[RFC2119] 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>.

[RFC8174] 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>.

[RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <https://www.rfc-editor.org/info/rfc8200>.

7.2. Informative References

[RFC2991] Thaler, D. and C. Hopps, "Multipath Issues in Unicast and Multicast Next-Hop Selection", RFC 2991, DOI 10.17487/RFC2991, November 2000, <https://www.rfc-editor.org/info/rfc2991>.

[RFC2992] Hopps, C., "Analysis of an Equal-Cost Multi-Path Algorithm", RFC 2992, DOI 10.17487/RFC2992, November 2000, <https://www.rfc-editor.org/info/rfc2992>.

[RFC3828] Larzon, L., Degermark, M., Pink, S., Jonsson, L., Ed., and G. Fairhurst, Ed., "The Lightweight User Datagram Protocol (UDP-Lite)", RFC 3828, DOI 10.17487/RFC3828, July 2004, <https://www.rfc-editor.org/info/rfc3828>.

[RFC4787] Audet, F., Ed. and C. Jennings, "Network Address Translation (NAT) Behavioral Requirements for Unicast UDP", BCP 127, RFC 4787, DOI 10.17487/RFC4787, January 2007, <https://www.rfc-editor.org/info/rfc4787>.

[RFC4987] Eddy, W., "TCP SYN Flooding Attacks and Common Mitigations", RFC 4987, DOI 10.17487/RFC4987, August 2007, <https://www.rfc-editor.org/info/rfc4987>.

[RFC6335] Cotton, M., Eggert, L., Touch, J., Westerlund, M., and S. Cheshire, "Internet Assigned Numbers Authority (IANA) Procedures for the Management of the Service Name and Transport Protocol Port Number Registry", BCP 165, RFC 6335, DOI 10.17487/RFC6335, August 2011, <https://www.rfc-editor.org/info/rfc6335>.

[RFC9293] Eddy, W., Ed., "Transmission Control Protocol (TCP)", STD 7, RFC 9293, DOI 10.17487/RFC9293, August 2022, <https://www.rfc-editor.org/info/rfc9293>.

Appendix A. Acknowledgments

The authors would like to thank all individuals who have provided valuable feedback and contributions during the development of this document.

Authors' Addresses

Zhaoyu Luo (editor)
CMCC
No. 58 Kunlunshan Road
Suzhou
215000
China
Email: luozhaoyu@cmss.chinamobile.com

Haishuang Yan
CMCC
No. 58 Kunlunshan Road
Suzhou
215000
China
Email: yanhaishuang@cmss.chinamobile.com