

IETF
Internet-Draft
Intended status: Informational
Expires: 11 August 2026

L. Chen
Zhongguancun Laboratory
D. Li
Tsinghua University
L. Liu
L. Qin
Zhongguancun Laboratory
7 February 2026

Benchmarking Methodology for Intra-domain and Inter-domain Source
Address Validation
draft-ietf-bmwg-savnet-sav-benchmarking-00

Abstract

This document defines methodologies for benchmarking the performance of intra-domain and inter-domain source address validation (SAV) mechanisms. SAV mechanisms are utilized to generate SAV rules to prevent source address spoofing, and have been implemented with many various designs in order to perform SAV in the corresponding scenarios. This document takes the approach of considering a SAV device to be a black box, defining the methodology in a manner that is agnostic to the mechanisms. This document provides a method for measuring the performance of existing and new SAV implementations.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 11 August 2026.

Copyright Notice

Copyright (c) 2026 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	3
1.1. Goal and Scope	3
1.2. Requirements Language	4
2. Terminology	4
3. Test Methodology	4
3.1. Test Setup	4
3.2. Network Topology and Device Configuration	5
4. SAV Performance Indicators	6
4.1. False Positive Rate	6
4.2. False Negative Rate	6
4.3. Protocol Convergence Time	6
4.4. Protocol Message Processing Throughput	6
4.5. Data Plane SAV Table Refreshing Rate	7
4.6. Data Plane Forwarding Rate	7
4.7. Resource Utilization	7
5. Benchmarking Tests	7
5.1. Intra-domain SAV	7
5.1.1. False Positive and False Negative Rates	7
5.1.2. Control Plane Performance	18
5.1.3. Data Plane Performance	20
5.2. Inter-domain SAV	22
5.2.1. False Positive and False Negative Rates	22
5.2.2. Control Plane Performance	38
5.2.3. Data Plane Performance	38
5.3. Resource Utilization	39
6. Reporting Format	39
7. IANA Considerations	39
8. Security Considerations	39
9. References	39
9.1. Normative References	40
9.2. Informative References	40
Acknowledgements	41
Authors' Addresses	41

1. Introduction

Source address validation (SAV) is significantly important to prevent source address spoofing. Operators are suggested to deploy different SAV mechanisms [RFC3704] [RFC8704] based on their deployment network environments. In addition, existing intra-domain (intra-AS) and inter-domain (inter-AS) SAV mechanisms have problems in operational overhead and SAV accuracy under various scenarios [intra-domain-ps] [inter-domain-ps]. Intra-domain and inter-domain SAVNET architectures [intra-domain-arch] [inter-domain-arch] are proposed to guide the design of new intra-domain and inter-domain SAV mechanisms to solve the problems. The benchmarking methodology defined in this document will help operators to get a more accurate idea of the SAV performance when their deployed devices enable SAV and will also help vendors to test the performance of SAV implementation for their devices.

This document provides generic methodologies for benchmarking SAV mechanism performance. To achieve the desired functionality, a SAV device may support multiple SAV mechanisms, allowing operators to enable those most suitable for their specific network environments. This document considers a SAV device to be a black box, regardless of the design and implementation. The tests defined in this document can be used to benchmark a SAV device for SAV accuracy (i.e., false positive and false negative rates), SAV protocol convergence performance, and control plane and data plane forwarding performance. These tests can be performed on a hardware router, a software router, a virtual machine (VM) instance, or a container instance, which runs as a SAV device. This document outlines methodologies for assessing SAV device performance and comparing various SAV mechanisms and implementations.

1.1. Goal and Scope

The benchmarking methodology outlined in this draft focuses on two objectives:

- * Assessing “which SAV mechanism performs best” over a set of well-defined scenarios.
- * Measuring the contribution of sub-systems to the overall SAV systems' performance (also known as “micro-benchmark”).

This benchmark evaluates the SAV performance of individual devices (e.g., hardware/software routers) by comparing different SAV mechanisms under specific network scenarios. The results help determine the appropriate SAV deployment for real-world network scenarios.

1.2. Requirements Language

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.

2. Terminology

SAV Control Plane: The SAV control plane consists of processes including gathering and communicating SAV-related information.

SAV Data Plane: The SAV data plane stores the SAV rules within a specific data structure and validates each incoming packet to determine whether to permit or discard it.

Host-facing Router: An edge router directly connected to a layer-2 host network.

Customer-facing Router: An edge router connected to a non-BGP customer network which includes routers and runs the routing protocol.

AS Border Router: An intra-domain router facing an external AS.

3. Test Methodology

3.1. Test Setup

The test setup in general is compliant with [RFC2544]. The Device Under Test (DUT) is connected to a Tester and other network devices to construct the network topology introduced in Section 5. The Tester is a traffic generator to generate network traffic with various source and destination addresses in order to emulate the spoofing or legitimate traffic. It is OPTIONAL to choose various proportions of traffic.

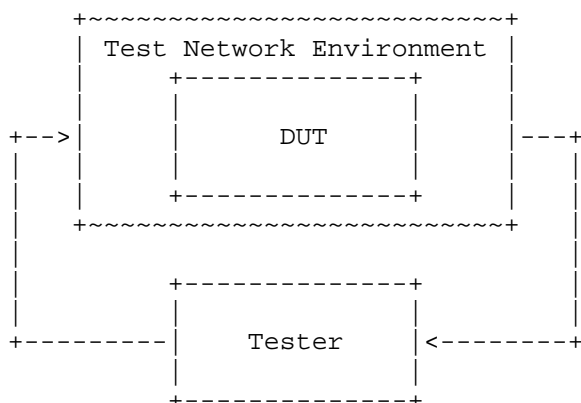


Figure 1: Test Setup.

Figure 1 illustrates the test configuration for the Device Under Test (DUT). Within the test network environment, the DUT can be interconnected with other devices to create a variety of test scenarios. The Tester may establish a direct connection with the DUT or link through intermediary devices. The nature of the connection between them is dictated by the benchmarking tests outlined in Section 5. Furthermore, the Tester has the capability to produce both spoofed and legitimate traffic to evaluate the SAV accuracy of the DUT in relevant scenarios, and it can also generate traffic at line rate to assess the data plane forwarding performance of the DUT. Additionally, the DUT is required to support logging functionalities to document all test outcomes.

3.2. Network Topology and Device Configuration

The positioning of the DUT within the network topology has an impact on SAV performance. Therefore, the benchmarking process MUST include evaluating the DUT at multiple locations across the network to ensure a comprehensive assessment.

The routing configurations of network devices may differ, and the resulting SAV rules depend on these settings. It is essential to clearly document the specific device configurations used during testing.

Furthermore, the role of each device, such as host-facing router, customer-facing router, or AS border router in an intra-domain network, SHOULD be clearly identified. In an inter-domain context, the business relationships between ASes MUST also be specified.

When evaluating data plane forwarding performance, the traffic generated by the Tester must be characterized by defined traffic rates, the ratio of spoofed to legitimate traffic, and the distribution of source addresses, as all of these factors can influence test results.

4. SAV Performance Indicators

This section lists key performance indicators (KPIs) of SAV for overall benchmarking tests. All KPIs SHOULD be measured in the benchmarking scenarios described in Section 5. Also, the KPIs SHOULD be measured from the result output of the DUT. The standard deviation for KPIs' testing results SHOULD be analyzed for each fixed test setup, which can help understand the stability of the DUT's performance. The data plane SAV table refreshing rate and data plane forwarding rate below SHOULD be tested using varying SAV table sizes for each fixed test setup, which can help measure DUT's sensibility to the SAV table size for these two KPIs.

4.1. False Positive Rate

The proportion of legitimate traffic which is determined to be spoofing traffic by the DUT across all the legitimate traffic, and this can reflect the SAV accuracy of the DUT.

4.2. False Negative Rate

The proportion of spoofing traffic which is determined to be legitimate traffic by the DUT across all the spoofing traffic, and this can reflect the SAV accuracy of the DUT.

4.3. Protocol Convergence Time

The control protocol convergence time represents the period during which the SAV control plane protocol converges to update the SAV rules when routing changes happen, and it is the time elapsed from the beginning of routing change to the completion of SAV rule update. This KPI can indicate the convergence performance of the SAV protocol.

4.4. Protocol Message Processing Throughput

The protocol message processing throughput measures the throughput of processing the packets for communicating SAV-related information on the control plane, and it can indicate the SAV control plane performance of the DUT.

4.5. Data Plane SAV Table Refreshing Rate

The data plane SAV table refreshing rate refers to the rate at which a DUT updates its SAV table with new SAV rules, and it can reflect the SAV data plane performance of the DUT.

4.6. Data Plane Forwarding Rate

The data plane forwarding rate measures the SAV data plane forwarding throughput for processing the data plane traffic, and it can indicate the SAV data plane performance of the DUT. It is suggested that measuring the data plane forwarding rate of DUT enabling and disabling SAV to see the proportion of decrease for the data plane forwarding rate. This can help analyze the efficiency for SAV data plane implementation of the DUT.

4.7. Resource Utilization

The resource utilization refers to the CPU and memory usage of the SAV processes within the DUT.

5. Benchmarking Tests

5.1. Intra-domain SAV

5.1.1. False Positive and False Negative Rates

***Objective*:** Evaluate the false positive rate and false negative rate of the DUT in processing both legitimate and spoofed traffic across various intra-domain network scenarios. These scenarios include SAV implementations for customer/host networks, Internet-facing networks, and aggregation-router-facing networks.

In the following, this document presents the test scenarios for evaluating intra-domain SAV performance on the DUT. Under each scenario, the generated spoofed traffic SHOULD include different types of forged source addresses, such as unused source addresses within the subnetwork, private network source addresses, internal-use-only source addresses of the subnetwork, and external source addresses. The ratios among these different types of forged source addresses SHOULD vary, since different SAV mechanisms may differ in their capability to block packets with forged source addresses of various types. Nevertheless, for all these types of spoofed traffic, the expected result is that the DUT SHOULD block them.

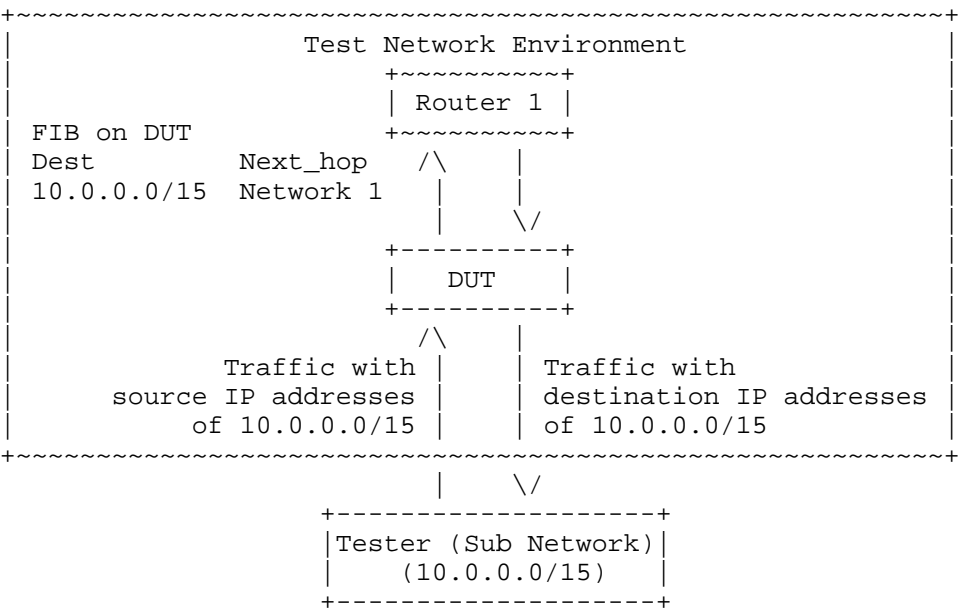


Figure 2: SAV for customer or host network in intra-domain symmetric routing scenario.

***SAV for Customer or Host Network*:** Figure 2 illustrates an intra-domain symmetric routing scenario in which SAV is deployed for a customer or host network. The DUT performs SAV as a customer/host-facing router and connects to Router 1 for Internet access. A sub network, which resides within the AS and uses the prefix 10.0.0.0/15, is connected to the DUT. The Tester emulates a sub network by advertising this prefix in the control plane and generating both spoofed and legitimate traffic in the data plane. In this setup, the Tester is configured so that inbound traffic destined for 10.0.0.0/15 arrives via the DUT. The DUT learns the route to 10.0.0.0/15 from the Tester, while the Tester sends outbound traffic with source addresses within 10.0.0.0/15 to the DUT, simulating a symmetric routing scenario between the two. The IP addresses used in this test case are optional; users may substitute them with other addresses, as applies equally to other test cases.

The ***procedure*** for testing SAV in this intra-domain symmetric routing scenario is as follows:

1. To verify whether the DUT can generate accurate SAV rules for customer or host network under symmetric routing conditions, construct a testbed as depicted in Figure 2. The Tester is connected to the DUT and acts as a sub network.

2. Configure the DUT and Router 1 to establish symmetric routing environment.
3. The Tester generates both legitimate traffic (with source addresses in 10.0.0.0/15) and spoofed traffic (with source addresses in 10.2.0.0/15) toward the DUT. The prefix 10.2.0.0/15 does not belong to the sub network and thus is not advertised by the Tester. The ratio of spoofed to legitimate traffic may vary, for example, from 1:9 to 9:1.

The **expected results** for this test case are that the DUT blocks spoofed traffic and allows legitimate traffic originating from the sub network.

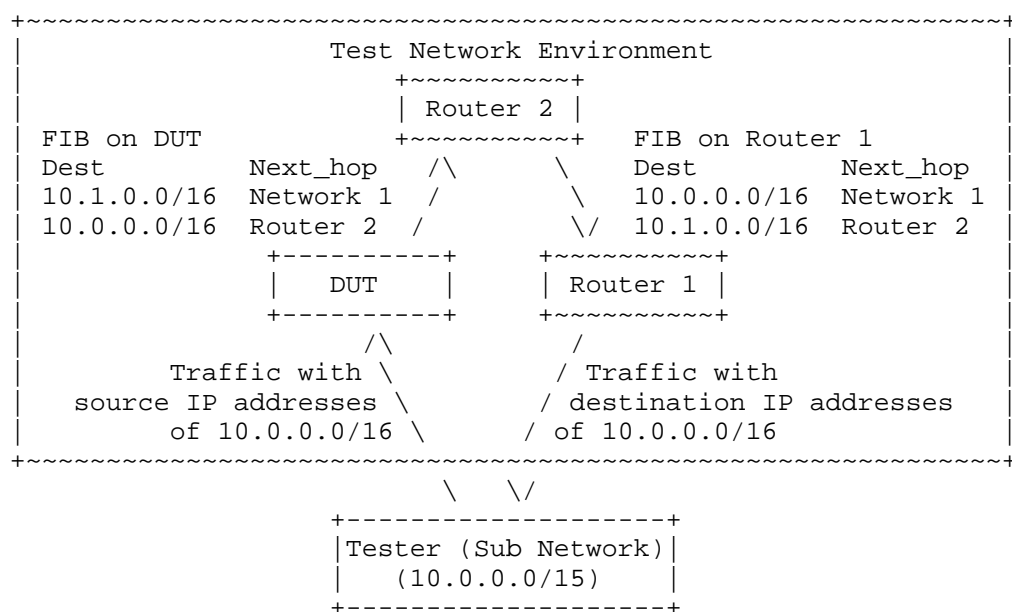


Figure 3: SAV for customer or host network in intra-domain asymmetric routing scenario.

SAV for Customer or Host Network: Figure 3 illustrates an intra-domain asymmetric routing scenario in which SAV is deployed for a customer or host network. The DUT performs SAV as a customer/host-facing router. A sub network, i.e., a customer/host network within the AS, is connected to both the DUT and Router 1, and uses the prefix 10.0.0.0/15. The Tester emulates a sub network and handles both its control plane and data plane functions. In this setup, the Tester is configured so that inbound traffic destined for 10.1.0.0/16 is received only from the DUT, while inbound traffic for 10.0.0.0/16

is received only from Router 1. The DUT learns the route to prefix 10.1.0.0/16 from the Tester, and Router 1 learns the route to 10.0.0.0/16 from the Tester. Both the DUT and Router 1 then advertise their respective learned prefixes to Router 2. Consequently, the DUT learns the route to 10.0.0.0/16 from Router 2, and Router 1 learns the route to 10.1.0.0/16 from Router 2. The Tester sends outbound traffic with source addresses in 10.0.0.0/16 to the DUT, simulating an asymmetric routing scenario between the Tester and the DUT.

The **procedure** for testing SAV in this intra-domain asymmetric routing scenario is as follows:

1. To determine whether the DUT can generate accurate SAV rules under asymmetric routing conditions, set up the test environment as shown in Figure 3. The Tester is connected to both the DUT and Router 1 and emulates the functions of a sub network.
2. Configure the DUT, Router 1, and Router 2 to establish the asymmetric routing scenario.
3. The Tester generates both spoofed traffic (using source addresses in 10.1.0.0/16) and legitimate traffic (using source addresses in 10.0.0.0/16) toward the DUT. The prefix 10.1.0.0/16 does not belong to the sub network and thus is not advertised by the Tester. The ratio of spoofed to legitimate traffic may vary, for example, from 1:9 to 9:1.

The **expected results** for this test case are that the DUT blocks spoofed traffic and permits legitimate traffic originating from the sub network.

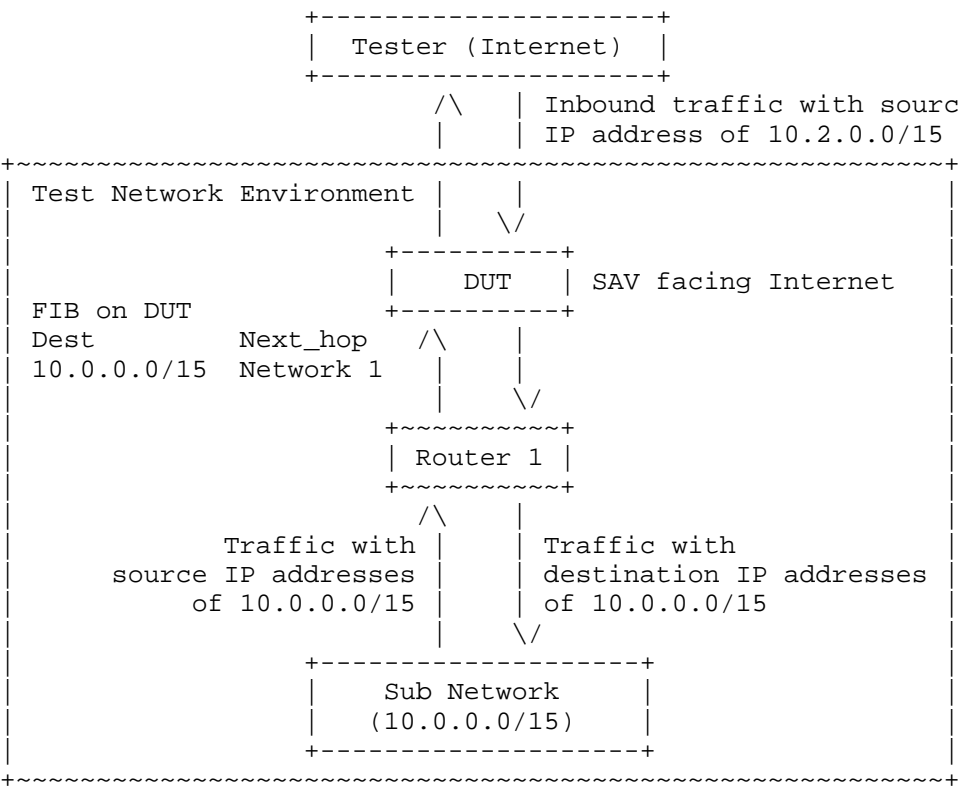


Figure 4: SAV for Internet-facing network in intra-domain symmetric routing scenario.

***SAV for Internet-facing Network*:** Figure 4 illustrates the test scenario for SAV in an Internet-facing network under intra-domain symmetric routing conditions. The network topology resembles that of Figure 2, with the key difference being the positioning of the DUT. In this case, the DUT is connected to Router 1 and the Internet, while the Tester emulates the Internet. The DUT performs SAV from an Internet-facing perspective, as opposed to a customer/host-facing role.

The *procedure* for testing SAV for an Internet-facing network in an intra-domain symmetric routing scenario is as follows:

1. To evaluate whether the DUT can generate accurate SAV rules for Internet-facing SAV under symmetric routing, set up the test environment as depicted in Figure 4. The Tester is connected to the DUT and emulates the Internet.

2. Configure the DUT and Router 1 to establish symmetric routing environment.
3. The Tester generates both spoofed traffic (using source addresses in 10.0.0.0/15) and legitimate traffic (using source addresses in 10.2.0.0/15) toward the DUT. The ratio of spoofed to legitimate traffic may vary, for example, from 1:9 to 9:1.

The *expected results* for this test case are that the DUT blocks spoofed traffic and allows legitimate traffic originating from the Internet.

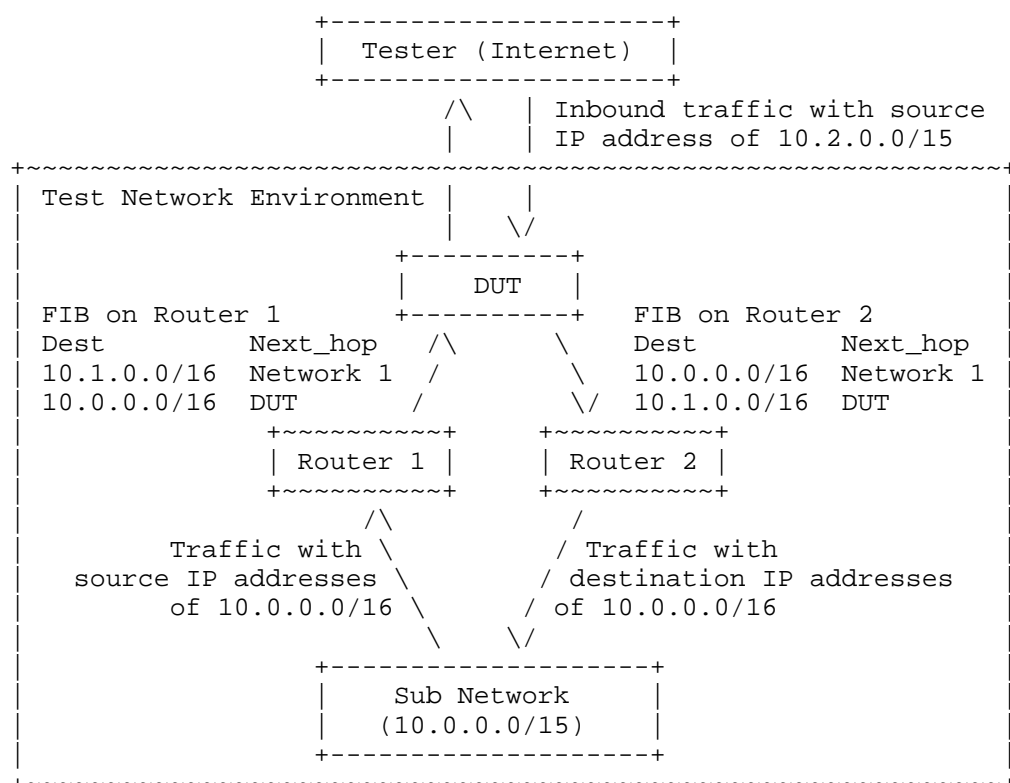


Figure 5: SAV for Internet-facing network in intra-domain asymmetric routing scenario.

SAV for Internet-facing Network: Figure 5 illustrates a test case for SAV in an Internet-facing network under intra-domain asymmetric routing conditions. The network topology is identical to that of Figure 3, with the key distinction being the placement of the DUT. In this scenario, the DUT is connected to Router 1 and Router 2

within the same AS, as well as to the Internet. The Tester emulates the Internet, and the DUT performs Internet-facing SAV rather than customer/host-network-facing SAV.

The *procedure* for testing SAV in this intra-domain asymmetric routing scenario is as follows:

1. To evaluate whether the DUT can generate accurate SAV rules for Internet-facing SAV under asymmetric routing, construct the test environment as shown in Figure 5. The Tester is connected to the DUT and emulates the Internet.
2. Configure the DUT, Router 1, and Router 2 to establish the asymmetric routing scenario.
3. The Tester generates both spoofed traffic (using source addresses in 10.0.0.0/15) and legitimate traffic (using source addresses in 10.2.0.0/15) toward the DUT. The ratio of spoofed to legitimate traffic may vary, for example, from 1:9 to 9:1.

The **expected results** for this test case are that the DUT blocks spoofed traffic and permits legitimate traffic originating from the Internet.

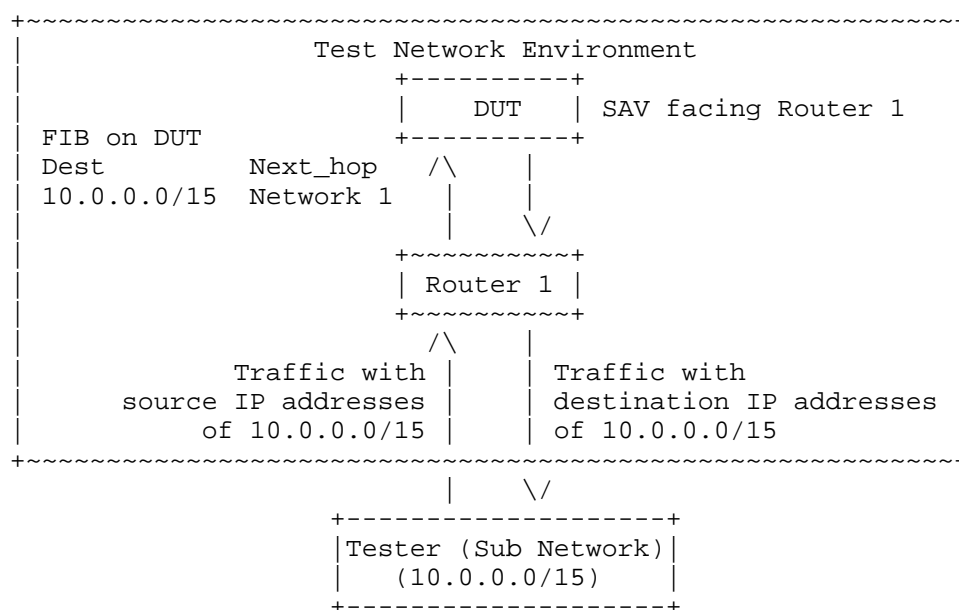


Figure 6: SAV for aggregation-router-facing network in intra-domain symmetric routing scenario.

***SAV for Aggregation-router-facing Network*:** Figure 6 depicts the test scenario for SAV in an aggregation-router-facing network under intra-domain symmetric routing conditions. The network topology in Figure 6 is identical to that of Figure 4. The Tester is connected to Router 1 to emulate a sub network, enabling evaluation of the DUT's false positive and false negative rates when facing Router 1.

The ***procedure*** for testing SAV in this aggregation-router-facing scenario is as follows:

1. To evaluate whether the DUT can generate accurate SAV rules for aggregation-router-facing SAV under symmetric routing, construct the test environment as shown in Figure 6. The Tester is connected to Router 1 and emulates a sub network.
2. Configure the DUT and Router 1 to establish symmetric routing environment.
3. The Tester generates both legitimate traffic (using source addresses in 10.1.0.0/15) and spoofed traffic (using source addresses in 10.2.0.0/15) toward Router 1. The prefix 10.2.0.0/15 does not belong to the sub network and thus is not advertised by the Tester. The ratio of spoofed to legitimate traffic may vary, for example, from 1:9 to 9:1.

The ***expected results*** for this test case are that the DUT blocks spoofed traffic and permits legitimate traffic originating from the direction of Router 1.

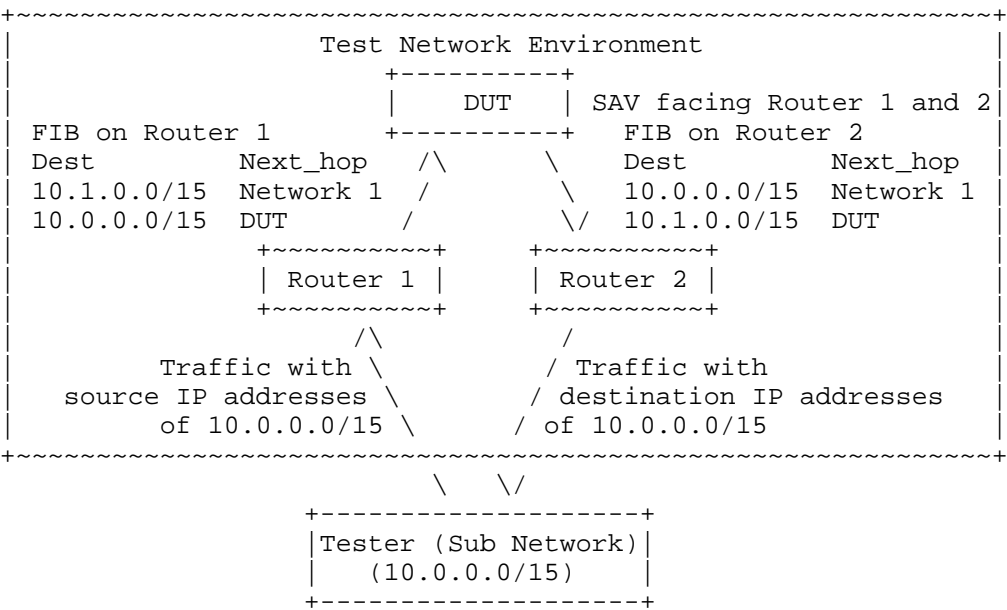


Figure 7: SAV for aggregation-router-facing network in intra-domain asymmetric routing scenario.

***SAV for Aggregation-router-facing Network*:** Figure 7 illustrates the test case for SAV in an aggregation-router-facing network under intra-domain asymmetric routing conditions. The network topology in Figure 7 is identical to that of Figure 5. The Tester is connected to both Router 1 and Router 2 to emulate a sub network, enabling evaluation of the DUT’s false positive and false negative rates when facing Router 1 and Router 2.

The ***procedure*** for testing SAV in this aggregation-router-facing asymmetric routing scenario is as follows:

1. To evaluate whether the DUT can generate accurate SAV rules under asymmetric routing conditions, construct the test environment as shown in Figure 7. The Tester is connected to Router 1 and Router 2 and emulates the functions of a sub network.
2. Configure the DUT, Router 1, and Router 2 to establish an asymmetric routing environment.

3. The Tester generates both spoofed traffic (using source addresses in 10.2.0.0/15) and legitimate traffic (using source addresses in 10.1.0.0/15) toward Router 1. The prefix 10.2.0.0/15 does not belong to the sub network and thus is not advertised by the Tester. The ratio of spoofed to legitimate traffic may vary, for example, from 1:9 to 9:1.

The **expected results** for this test case are that the DUT blocks spoofed traffic and permits legitimate traffic originating from the direction of Router 1 and Router 2.

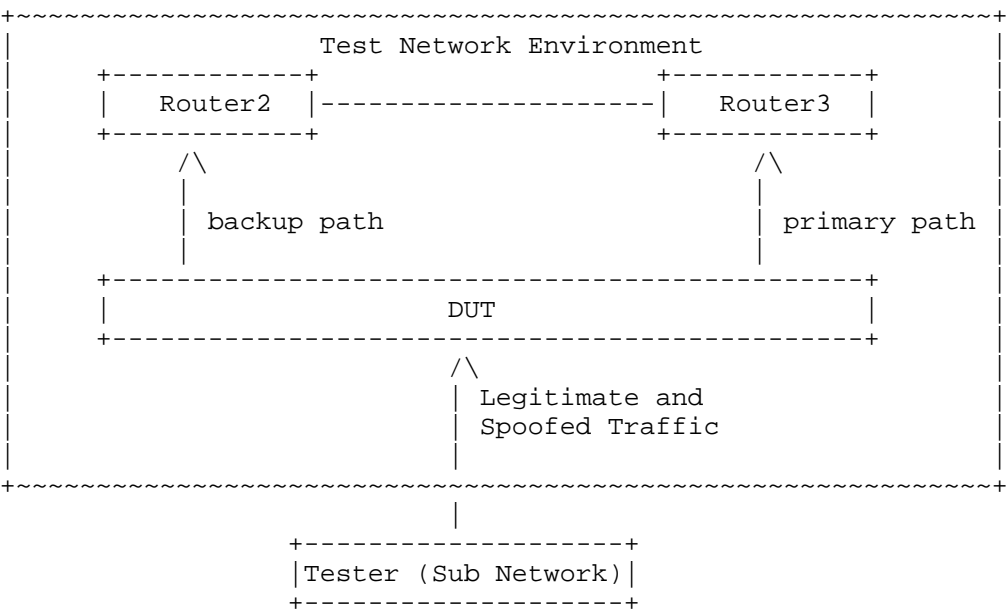


Figure 8: Intra-domain SAV under Fast Reroute (FRR) scenario.

SAV under Fast Reroute (FRR) Scenario: Fast Reroute (FRR) mechanisms such as Loop-Free Alternates (LFA) or Topology-Independent Loop-Free Alternates (TI-LFA) provide sub-second restoration of traffic forwarding after link or node failures. During FRR activation, temporary forwarding changes may occur before the control plane converges, potentially impacting SAV rule consistency and causing transient false positives or false negatives.

The **procedure** for testing SAV under FRR scenario is as follows:

1. Configure the DUT and adjacent routers with FRR protection for the primary link (Router3DUT).

2. The Tester continuously sends legitimate and spoofed traffic toward the protected prefix.
3. Trigger a link failure between Router3 and the DUT, causing FRR switchover to Router2.
4. Measure false positive and false negative rates during the switchover and after reconvergence.
5. Restore the primary link and verify that SAV rules revert correctly.

The **expected results** for this test case are that the DUT should maintain correct SAV behavior throughout FRR activation and recovery. False positive and false negative rates *SHOULD* remain minimal during FRR events, and SAV rules *SHOULD* update promptly to reflect restored routing.

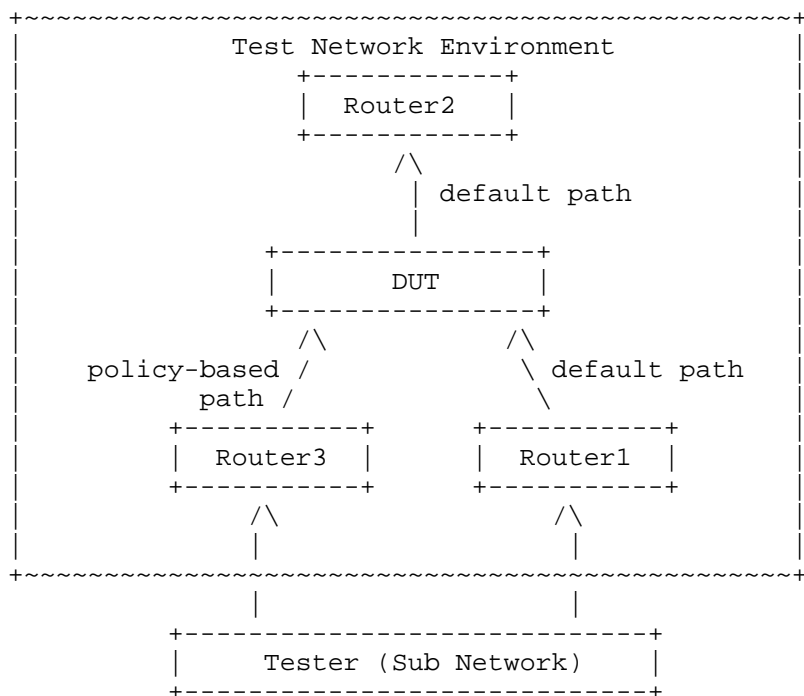


Figure 9: Intra-domain SAV under Policy-based Routing (PBR) scenario.

SAV under Policy-based Routing (PBR) Scenario: Policy-based Routing (PBR) enables forwarding decisions based on user-defined match conditions (e.g., source prefix, DSCP, or interface) instead of the

standard routing table. Such policies can create asymmetric paths that challenge the SAV mechanism if rules are derived solely from RIB or FIB information.

The **procedure** for testing SAV under PBR scenario is as follows:

1. Configure PBR on the DUT to forward traffic matching a specific source prefix (e.g., 10.1.0.0/16) to Router3, while other traffic follows the default path to Router1.
2. The Tester sends both legitimate and spoofed traffic that matches and does not match the PBR policy.
3. Measure the false positive and false negative rates for both traffic types.
4. Dynamically modify or remove the PBR policy and observe SAV rule adaptation.

The **expected results** for this test case are that the DUT SHOULD continue to correctly filter spoofed packets and permit legitimate packets under PBR scenario. SAV rules MUST adapt to policy-based forwarding paths without producing misclassification.

5.1.2. Control Plane Performance

Objective: Measure the control plane performance of the DUT, including both protocol convergence performance and protocol message processing performance in response to route changes caused by network failures or operator configurations. Protocol convergence performance is quantified by the convergence time, defined as the duration from the onset of a routing change until the completion of the corresponding SAV rule update. Protocol message processing performance is measured by the processing throughput, represented by the total size of protocol messages processed per second.

Note that the tests for control plane performance of the DUT which performs intra-domain SAV are OPTIONAL. Only DUT which implements the SAV mechanism using an explicit control-plane communication protocol, such as SAV-specific information communication mechanism proposed in [intra-domain-arch] SHOULD be tested on its control plane performance.

```
+~~~~~+      +-----+      +-----+
| Emulated Topology |-----|   Tester   |<----->|   DUT   |
+~~~~~+      +-----+      +-----+
```

Figure 10: Test setup for protocol convergence performance measurement.

***Protocol Convergence Performance*:** Figure 10 illustrates the test setup for measuring protocol convergence performance. The convergence process of the DUT, during which SAV rules are updated, is triggered by route changes resulting from network failures or operator configurations. In Figure 10, the Tester is directly connected to the DUT and simulates these route changes by adding or withdrawing prefixes to initiate the DUT's convergence procedure.

The ***procedure*** for testing protocol convergence performance is as follows:

1. To measure the protocol convergence time of the DUT, set up the test environment as depicted in Figure 10, with the Tester directly connected to the DUT.
2. The Tester withdraws a specified percentage of the total prefixes supported by the DUT, for example, 10%, 20%, up to 100%.
3. The protocol convergence time is calculated based on DUT logs that record the start and completion times of the convergence process.

Please note that for IGP, proportional prefix withdrawal can be achieved by selectively shutting down interfaces. For instance, if the Tester is connected to ten emulated devices through ten interfaces, each advertising a prefix, withdrawing 10% of prefixes can be accomplished by randomly disabling one interface. Similarly, 20% withdrawal corresponds to shutting down two interfaces, and so forth. This is one suggested method, and other approaches that achieve the same effect should be also acceptable.

The protocol convergence time, defined as the duration required for the DUT to complete the convergence process, should be measured from the moment the last "hello" message is received from the emulated device on the disabled interface until SAV rule generation is finalized. To ensure accuracy, the DUT should log the timestamp of the last hello message received and the timestamp when SAV rule updates are complete. The convergence time is the difference between these two timestamps.

It is recommended that if the emulated device sends a "goodbye hello" message during interface shutdown, using the receipt time of this message, rather than the last standard hello, as the starting point will provide a more precise measurement, as advised in [RFC4061].

***Protocol Message Processing Performance*:** The test for protocol message processing performance uses the same setup illustrated in Figure 10. This performance metric evaluates the protocol message processing throughput, the rate at which the DUT processes protocol messages. The Tester varies the sending rate of protocol messages, ranging from 10% to 100% of the total link capacity between the Tester and the DUT. The DUT records both the total size of processed protocol messages and the corresponding processing time.

The ***procedure*** for testing protocol message processing performance is as follows:

1. To measure the protocol message processing throughput of the DUT, set up the test environment as shown in Figure 10, with the Tester directly connected to the DUT.
2. The Tester sends protocol messages at varying rates, such as 10%, 20%, up to 100%, of the total link capacity between the Tester and the DUT.
3. The protocol message processing throughput is calculated based on DUT logs that record the total size of processed protocol messages and the total processing time.

To compute the protocol message processing throughput, the DUT logs MUST include the total size of the protocol messages processed and the total time taken for processing. The throughput is then derived by dividing the total message size by the total processing time.

5.1.3. Data Plane Performance

***Objective*:** Evaluate the data plane performance of the DUT, including both data plane SAV table refresh performance and data plane forwarding performance. Data plane SAV table refresh performance is quantified by the refresh rate, which indicates how quickly the DUT updates its SAV table with new SAV rules. Data plane forwarding performance is measured by the forwarding rate, defined as the total size of packets forwarded by the DUT per second.

***Data Plane SAV Table Refreshing Performance*:** The evaluation of data plane SAV table refresh performance uses the same test setup shown in Figure 10. This metric measures the rate at which the DUT refreshes its SAV table with new SAV rules. The Tester varies the transmission rate of protocol messages, from 10% to 100% of the total link capacity between the Tester and the DUT, to influence the proportion of updated SAV rules and corresponding SAV table entries. The DUT records the total number of updated SAV table entries and the time taken to complete the refresh process.

The **procedure** for testing data plane SAV table refresh performance is as follows:

1. To measure the data plane SAV table refreshing rate of the DUT, set up the test environment as depicted in Figure 10, with the Tester directly connected to the DUT.
2. The Tester sends protocol messages at varying percentages of the total link capacity, for example, 10%, 20%, up to 100%.
3. The data plane SAV table refreshing rate is calculated based on DUT logs that record the total number of updated SAV table entries and the total refresh time.

To compute the refresh rate, the DUT logs MUST capture the total number of updated SAV table entries and the total time required for refreshing. The refresh rate is then derived by dividing the total number of updated entries by the total refresh time.

Data Plane Forwarding Performance: The evaluation of data plane forwarding performance uses the same test setup shown in Figure 10. The Tester transmits a mixture of spoofed and legitimate traffic at a rate matching the total link capacity between the Tester and the DUT, while the DUT maintains a fully populated SAV table. The ratio of spoofed to legitimate traffic can be varied within a range, for example, from 1:9 to 9:1. The DUT records the total size of forwarded packets and the total duration of the forwarding process.

The procedure for testing data plane forwarding performance is as follows:

1. To measure the data plane forwarding rate of the DUT, set up the test environment as depicted in Figure 10, with the Tester directly connected to the DUT.
2. The Tester sends a mix of spoofed and legitimate traffic to the DUT at the full link capacity between the Tester and the DUT. The ratio of spoofed to legitimate traffic may vary, for example, from 1:9 to 9:1.
3. The data plane forwarding rate is calculated based on DUT logs that record the total size of forwarded traffic and the total forwarding time.

To compute the forwarding rate, the DUT logs must include the total size of forwarded traffic and the total time taken for forwarding. The forwarding rate is then derived by dividing the total traffic size by the total forwarding time.

5.2. Inter-domain SAV

5.2.1. False Positive and False Negative Rates

***Objective*:** Measure the false positive rate and false negative rate of the DUT when processing legitimate and spoofed traffic across multiple inter-domain network scenarios, including SAV implementations for both customer-facing ASes and provider-/peer-facing ASes.

In the following, this document presents the test scenarios for evaluating inter-domain SAV performance on the DUT. Under each scenario, the generated spoofed traffic SHOULD include different types of forged source addresses, such as source addresses belonging to the local AS but not announced to external networks, private network source addresses, source addresses belonging to other ASes, and unallocated (unused) source addresses. The ratios among these different types of forged source addresses SHOULD vary, since different inter-domain SAV mechanisms may differ in their capability to block packets with forged source addresses of various origins. Nevertheless, for all these types of spoofed traffic, the expected result is that the DUT SHOULD block them.

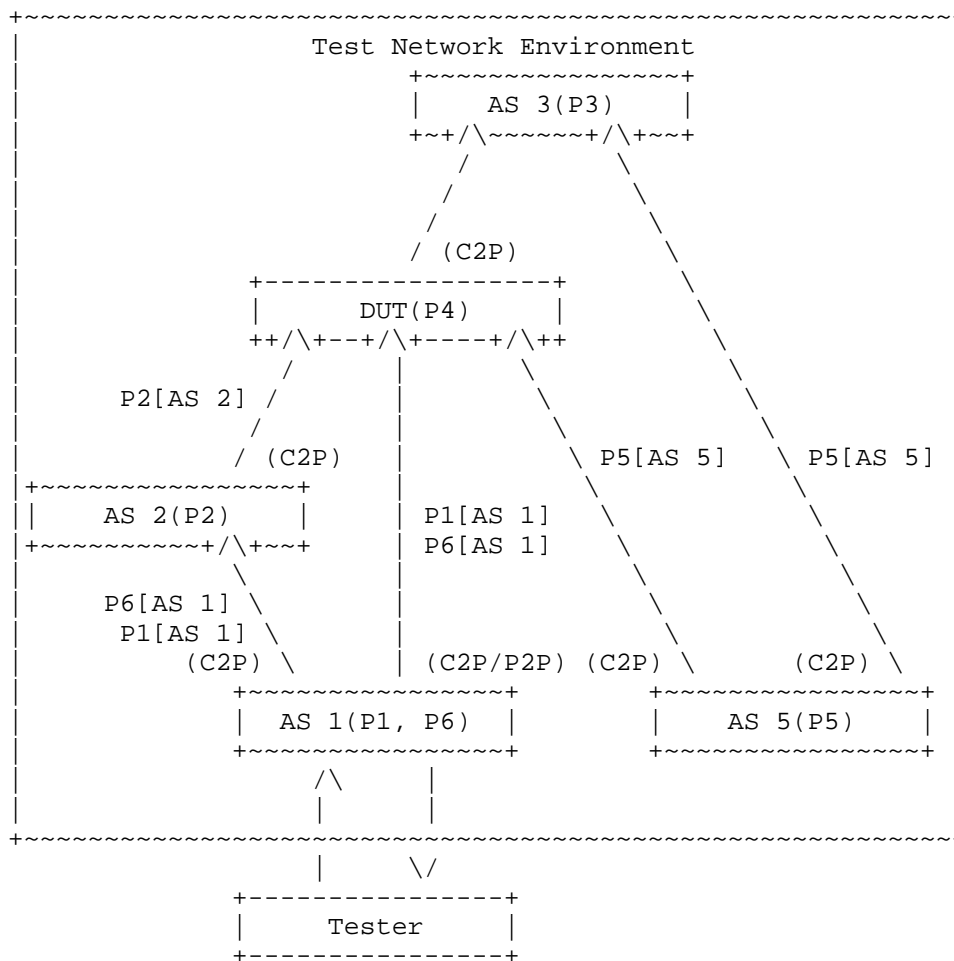


Figure 11: SAV for customer-facing ASes in inter-domain symmetric routing scenario.

***SAV for Customer-facing ASes*:** Figure 11 presents a test case for SAV in customer-facing ASes under an inter-domain symmetric routing scenario. In this setup, AS 1, AS 2, AS 3, the DUT, and AS 5 form the test network environment, with the DUT performing SAV at the AS level. AS 1 is a customer of both AS 2 and the DUT; AS 2 is a customer of the DUT, which in turn is a customer of AS 3; and AS 5 is a customer of both AS 3 and the DUT. AS 1 advertises prefixes P1 and P6 to AS 2 and the DUT, respectively. AS 2 then propagates routes for P1 and P6 to the DUT, enabling the DUT to learn these prefixes from both AS 1 and AS 2. In this test, the legitimate path for traffic with source addresses in P1 and destination addresses in P4 is AS 1->AS 2->DUT->AS 4. The Tester is connected to AS 1 to evaluate the DUT's SAV performance for customer-facing ASes.

The ***procedure*** for testing SAV in this scenario is as follows:

1. To evaluate whether the DUT can generate accurate SAV rules for customer-facing ASes under symmetric inter-domain routing, construct the test environment as shown in Figure 11. The Tester is connected to AS 1 and generates test traffic toward the DUT.
2. Configure AS 1, AS 2, AS 3, the DUT, and AS 5 to establish symmetric routing environment.
3. The Tester sends both legitimate traffic (with source addresses in P1 and destination addresses in P4) and spoofed traffic (with source addresses in P5 and destination addresses in P4) to the DUT via AS 2. The ratio of spoofed to legitimate traffic may vary, for example, from 1:9 to 9:1.

The ***expected results*** for this test case are that the DUT blocks spoofed traffic and permits legitimate traffic received from the direction of AS 2.

Note that the DUT may also be placed at AS 1 or AS 2 in Figure 11 to evaluate its false positive and false negative rates using the same procedure. In these configurations, the DUT is expected to effectively block spoofed traffic.

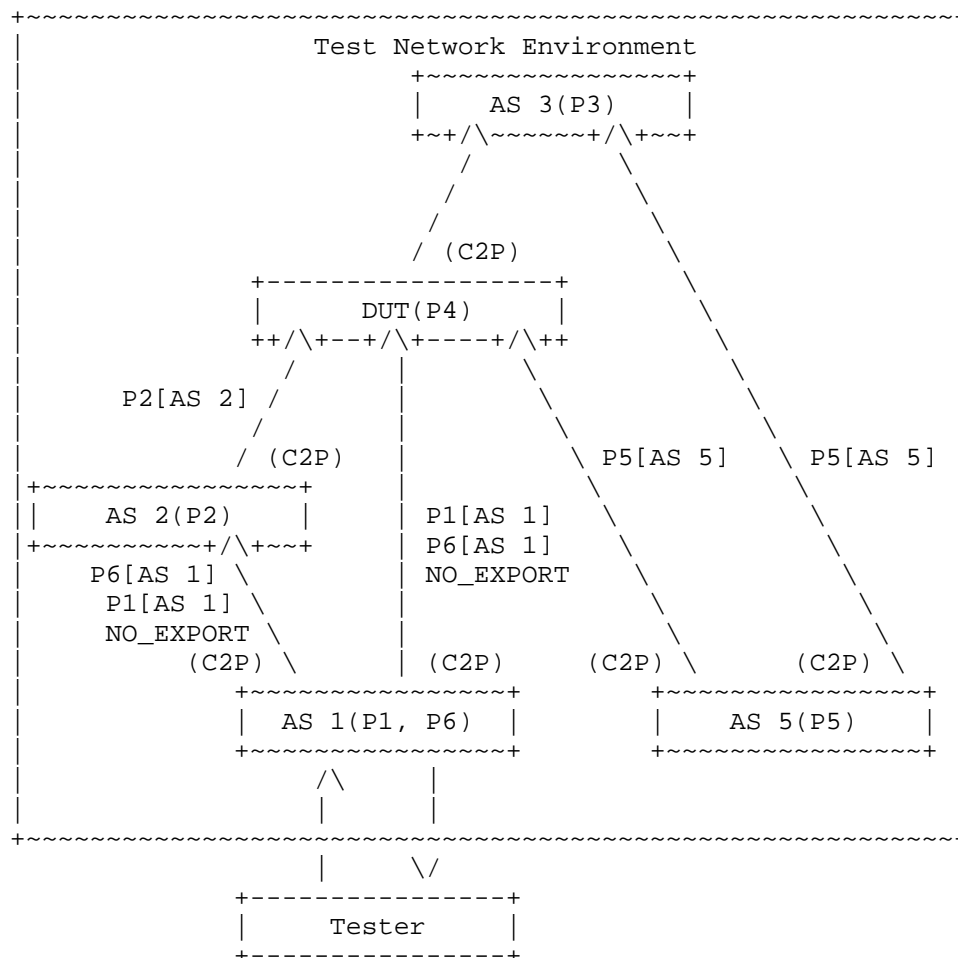


Figure 12: SAV for customer-facing ASes in inter-domain asymmetric routing scenario caused by NO EXPORT.

SAV for Customer-facing ASes: Figure 12 presents a test case for SAV in customer-facing ASes under an inter-domain asymmetric routing scenario induced by NO_EXPORT community configuration. In this setup, AS 1, AS 2, AS 3, the DUT, and AS 5 form the test network, with the DUT performing SAV at the AS level. AS 1 is a customer of both AS 2 and the DUT; AS 2 is a customer of the DUT, which is itself a customer of AS 3; and AS 5 is a customer of both AS 3 and the DUT. AS 1 advertises prefix P1 to AS 2 with the NO_EXPORT community attribute, preventing AS 2 from propagating the route for P1 to the DUT. Similarly, AS 1 advertises prefix P6 to the DUT with the NO_EXPORT attribute, preventing the DUT from propagating this route


to AS 3. As a result, the DUT learns the route for prefix P1 only from AS 1. The legitimate path for traffic with source addresses in P1 and destination addresses in P4 is AS 1->AS 2->DUT. The Tester is connected to AS 1 to evaluate the DUT's SAV performance for customer-facing ASes.

The **procedure** for testing SAV in this asymmetric routing scenario is as follows:

1. To evaluate whether the DUT can generate accurate SAV rules under NO_EXPORT-induced asymmetric routing, construct the test environment as shown in Figure 12. The Tester is connected to AS 1 and generates test traffic toward the DUT.
2. Configure AS 1, AS 2, AS 3, the DUT, and AS 5 to establish the asymmetric routing scenario.
3. The Tester sends both legitimate traffic (with source addresses in P1 and destination addresses in P4) and spoofed traffic (with source addresses in P5 and destination addresses in P4) to the DUT via AS 2. The ratio of spoofed to legitimate traffic may vary—for example, from 1:9 to 9:1.

The **expected results** for this test case are that the DUT blocks spoofed traffic and permits legitimate traffic received from the direction of AS 2.

Note that the DUT may also be placed at AS 1 or AS 2 in Figure 12 to evaluate its false positive and false negative rates using the same procedure. In these configurations, the DUT is expected to effectively block spoofed traffic.



SAV for Customer-facing ASes: Figure 13 presents a test case for SAV in customer-facing ASes under a Direct Server Return (DSR) scenario. In this setup, AS 1, AS 2, AS 3, the DUT, and AS 5 form the test network, with the DUT performing SAV at the AS level. AS 1 is a customer of both AS 2 and the DUT; AS 2 is a customer of the DUT, which is itself a customer of AS 3; and AS 5 is a customer of both AS 3 and the DUT. When users in AS 2 send requests to an anycast destination IP, the forwarding path is AS 2->DUT->AS 3.


Anycast servers in AS 3 receive the requests and tunnel them to edge servers in AS 1. The edge servers then return content to the users with source addresses in prefix P3. If the reverse forwarding path is AS 1->DUT->AS 2, the Tester sends traffic with source addresses in P3 and destination addresses in P2 along the path AS 1->DUT->AS 2. Alternatively, if the reverse forwarding path is AS 1->AS 2, the Tester sends traffic with source addresses in P3 and destination addresses in P2 along the path AS 1->AS 2. In this case, AS 2 may serve as the DUT.

The **procedure** for testing SAV in this DSR scenario is as follows:

1. To evaluate whether the DUT can generate accurate SAV rules under DSR conditions, construct the test environment as shown in Figure 13. The Tester is connected to AS 1 and generates test traffic toward the DUT.
2. Configure AS 1, AS 2, AS 3, the DUT, and AS 5 to establish the DSR scenario.
3. The Tester sends legitimate traffic (with source addresses in P3 and destination addresses in P2) to AS 2 via the DUT.

The **expected results** for this test case are that the DUT permits legitimate traffic with source addresses in P3 received from the direction of AS 1.

Note that the DUT may also be placed at AS 1 or AS 2 in Figure 13 to evaluate its false positive and false negative rates using the same procedure. In these configurations, the DUT is expected to effectively block spoofed traffic.



P1' is the spoofed source prefix P1 by the attacker which is inside of AS 2 or connected to AS 2 through other ASes.

Figure 14: SAV for customer-facing ASes in the scenario of reflection attacks.


***SAV for Customer-facing ASes*:** Figure 14 illustrates a test case for SAV in customer-facing ASes under a reflection attack scenario. In this scenario, a reflection attack using source address spoofing occurs within the DUT's customer cone. The attacker spoofs the victim's IP address (P1) and sends requests to server IP addresses (P5) that are configured to respond to such requests. The Tester emulates the attacker by performing source address spoofing. The arrows in Figure 14 indicate the business relationships between ASes: AS 3 serves as the provider for both the DUT and AS 5, while the DUT acts as the provider for AS 1, AS 2, and AS 5. Additionally, AS 2 is the provider for AS 1.

The **procedure** for testing SAV under reflection attack conditions is as follows:

1. To evaluate whether the DUT can generate accurate SAV rules in a reflection attack scenario, construct the test environment as shown in Figure 14. The Tester is connected to AS 2 and generates test traffic toward the DUT.
2. Configure AS 1, AS 2, AS 3, the DUT, and AS 5 to simulate the reflection attack scenario.
3. The Tester sends spoofed traffic (with source addresses in P1 and destination addresses in P5) toward AS 5 via the DUT.

The **expected results** for this test case are that the DUT blocks spoofed traffic with source addresses in P1 received from the direction of AS 2.

Note that the DUT may also be placed at AS 1 or AS 2 in Figure 14 to evaluate its false positive and false negative rates using the same procedure. In these configurations, the DUT is expected to effectively block spoofed traffic.



P5' is the spoofed source prefix P5 by the attacker which is inside of AS 2 or connected to AS 2 through other ASes.

Figure 15: SAV for customer-facing ASes in the scenario of direct attacks.


***SAV for Customer-facing ASes*:** Figure 15 presents a test case for SAV in customer-facing ASes under a direct attack scenario. In this scenario, a direct attack using source address spoofing occurs within the DUT's customer cone. The attacker spoofs a source address (P5) and directly targets the victim's IP address (P1), aiming to overwhelm its network resources. The Tester emulates the attacker by performing source address spoofing. The arrows in Figure 15 indicate the business relationships between ASes: AS 3 serves as the provider for both the DUT and AS 5, while the DUT acts as the provider for AS 1, AS 2, and AS 5. Additionally, AS 2 is the provider for AS 1.

The *procedure* for testing SAV under direct attack conditions is as follows:

1. To evaluate whether the DUT can generate accurate SAV rules in a direct attack scenario, construct the test environment as shown in Figure 15. The Tester is connected to AS 2 and generates test traffic toward the DUT.
2. Configure AS 1, AS 2, AS 3, the DUT, and AS 5 to simulate the direct attack scenario.
3. The Tester sends spoofed traffic (with source addresses in P5 and destination addresses in P1) toward AS 1 via the DUT.

The **expected results** for this test case are that the DUT blocks spoofed traffic with source addresses in P5 received from the direction of AS 2.

Note that DUT may also be placed at AS 1 or AS 2 in Figure 15 to evaluate its false positive and false negative rates using the same procedure. In these configurations, the DUT is expected to effectively block spoofed traffic.



P1' is the spoofed source prefix P1 by the attacker which is inside of AS 3 or connected to AS 3 through other ASes.

Figure 16: SAV for provider-facing ASes in the scenario of reflection attacks.

***SAV for Provider/Peer-facing ASes*:** Figure 16 illustrates a test case for SAV in provider/peer-facing ASes under a reflection attack scenario. In this scenario, the attacker spoofs the victim's IP address (P1) and sends requests to server IP addresses (P2) that are configured to respond. The Tester emulates the attacker by performing source address spoofing. The servers then send


overwhelming responses to the victim, exhausting its network resources. The arrows in Figure 16 represent the business relationships between ASes: AS 3 acts as either a provider or a lateral peer of the DUT and is the provider for AS 5, while the DUT serves as the provider for AS 1, AS 2, and AS 5. Additionally, AS 2 is the provider for AS 1.

The **procedure** for testing SAV under reflection attack conditions is as follows:

1. To evaluate whether the DUT can generate accurate SAV rules for provider/peer-facing ASes in a reflection attack scenario, construct the test environment as shown in Figure 16. The Tester is connected to AS 3 and generates test traffic toward the DUT.
2. Configure AS 1, AS 2, AS 3, the DUT, and AS 5 to simulate the reflection attack scenario.
3. The Tester sends spoofed traffic (with source addresses in P1 and destination addresses in P2) toward AS 2 via AS 3 and the DUT.

The **expected results** for this test case are that the DUT blocks spoofed traffic with source addresses in P1 received from the direction of AS 3.

Note that the DUT may also be placed at AS 1 or AS 2 in Figure 16 to evaluate its false positive and false negative rates using the same procedure. In these configurations, the DUT is expected to effectively block spoofed traffic.



P2' is the spoofed source prefix P2 by the attacker which is inside of AS 3 or connected to AS 3 through other ASes.

Figure 17: SAV for provider-facing ASes in the scenario of direct attacks.

Figure 17 presents a test case for SAV in provider-facing ASes under a direct attack scenario. In this scenario, the attacker spoofs a source address (P2) and directly targets the victim's IP address (P1), overwhelming its network resources. The arrows in Figure 17 represent the business relationships between ASes: AS 3 acts as either a provider or a lateral peer of the DUT and is the provider for AS 5, while the DUT serves as the provider for AS 1, AS 2, and AS 5. Additionally, AS 2 is the provider for AS 1.

The procedure for testing SAV under direct attack conditions is as follows:

1. To evaluate whether the DUT can generate accurate SAV rules for provider-facing ASes in a direct attack scenario, construct the test environment as shown in Figure 17. The Tester is connected to AS 3 and generates test traffic toward the DUT.
2. Configure AS 1, AS 2, AS 3, the DUT, and AS 5 to simulate the direct attack scenario.
3. The Tester sends spoofed traffic (with source addresses in P2 and destination addresses in P1) toward AS 1 via AS 3 and the DUT.

The **expected results** for this test case are that the DUT blocks spoofed traffic with source addresses in P2 received from the direction of AS 3.

Note that the DUT may also be placed at AS 1 or AS 2 in Figure 17 to evaluate its false positive and false negative rates using the same procedure. In these configurations, the DUT is expected to effectively block spoofed traffic.

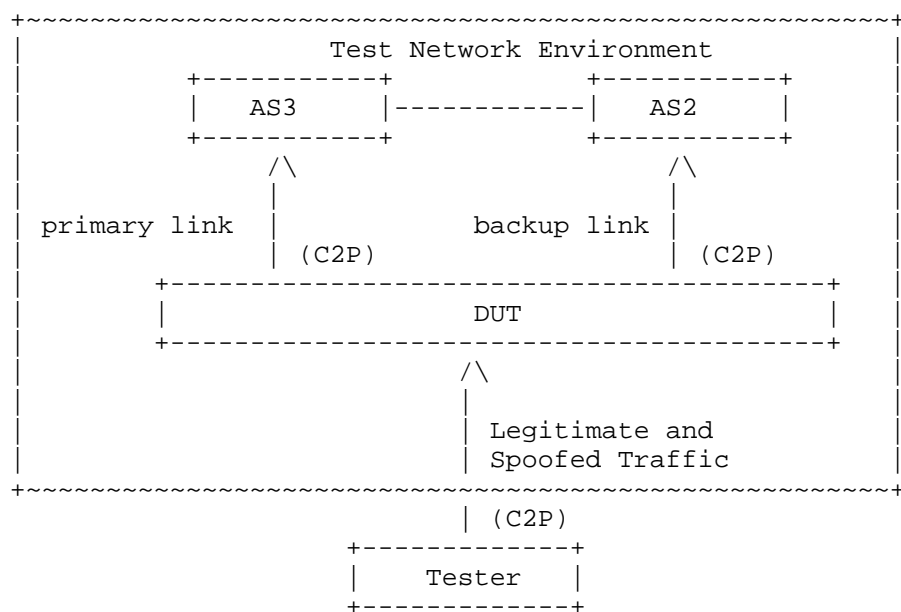


Figure 18: Inter-domain SAV under FRR scenario.

SAV under FRR Scenario: Inter-domain Fast Reroute (FRR) mechanisms, such as BGP Prefix Independent Convergence (PIC) or MPLS-based FRR, allow rapid failover between ASes after a link or node failure. These events may temporarily desynchronize routing information and SAV rules.

The *procedure* for testing SAV under FRR scenario is as follows:

1. Configure FRR or BGP PIC on the DUT for inter-AS links to AS3 (primary) and AS2 (backup).
2. Continuously send legitimate and spoofed traffic from AS1 toward DUT.
3. Trigger a failure on the AS3DUT link to activate the FRR path via AS2.
4. Measure false positive and false negative rates during and after switchover.
5. Restore the AS3 link and verify SAV table consistency.

The *expected results* for this test case are that the DUT MUST maintain consistent SAV filtering during FRR events. Transient topology changes SHOULD NOT lead to acceptance of spoofed traffic or unnecessary blocking of legitimate packets.

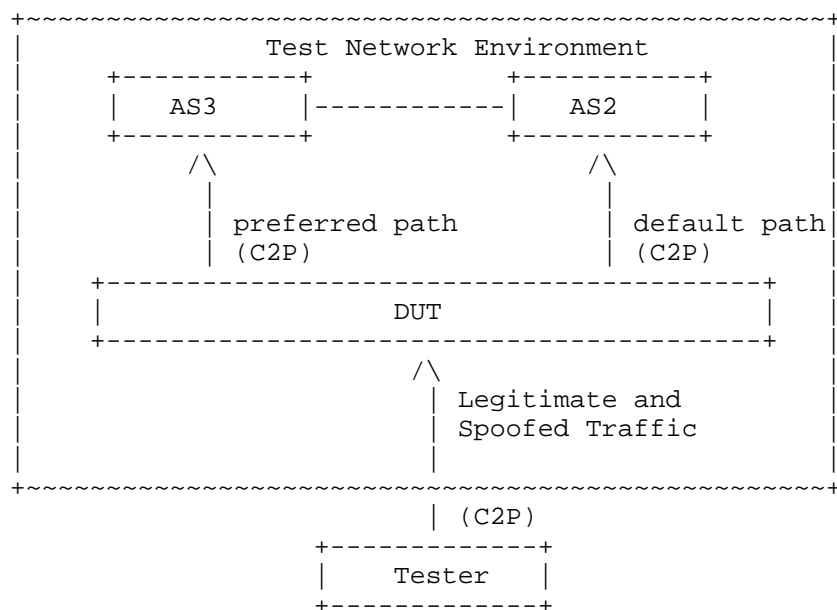


Figure 19: Inter-domain SAV under PBR scenario.

***SAV under PBR Scenario*:** In inter-domain environments, routing policies such as local preference, route maps, or communities may alter path selection independently of shortest-path routing. Such policy-driven forwarding can affect how the SAV rules are derived and applied.

The ***procedure*** for testing SAV under PBR scenario is as follows:

1. Configure a routing policy on the DUT (e.g., set local preference) to prefer AS3 for specific prefixes while maintaining AS2 as an alternative path.
2. Generate legitimate and spoofed traffic from AS1 matching both policy-affected and unaffected prefixes.
3. Observe SAV filtering behavior before and after policy changes.
4. Modify the routing policy dynamically and measure false positive and false negative rates.

The ***expected results*** for this test case are that the DUT SHOULD maintain correct SAV filtering regardless of routing policy changes. Legitimate traffic rerouted by policy MUST NOT be dropped, and spoofed traffic MUST NOT be forwarded during or after policy updates.

5.2.2. Control Plane Performance

The test setup, procedure, and metrics for evaluating protocol convergence performance and protocol message processing performance can refer to Section 5.1.2. Note that the tests for control plane performance of the DUT which performs inter-domain SAV are OPTIONAL. Only DUT which implements the SAV mechanism using an explicit control-plane communication protocol, such as SAV-specific information communication mechanism proposed in [inter-domain-arch] SHOULD be tested on its control plane performance.

5.2.3. Data Plane Performance

The test setup, procedure, and metrics for evaluating data plane SAV table refresh performance and data plane forwarding performance can refer to Section 5.1.3.

5.3. Resource Utilization

When evaluating the DUT for both intra-domain (Section 5.1) and inter-domain SAV (Section 5.2) functionality, CPU utilization (for both control and data planes) and memory utilization (for both control and data planes) MUST be recorded. These metrics SHOULD be collected separately per plane to facilitate granular performance analysis.

6. Reporting Format

Each test follows a reporting format comprising both global, standardized components and individual elements specific to each test. The following parameters for test configuration and SAV mechanism settings MUST be documented in the test report.

Test Configuration Parameters:

1. Test device hardware and software versions
2. Network topology
3. Test traffic attributes
4. System configuration (e.g., physical or virtual machine, CPU, memory, caches, operating system, interface capacity)
5. Device configuration (e.g., symmetric routing, NO_EXPORT)
6. SAV mechanism

7. IANA Considerations

This document has no IANA actions.

8. Security Considerations

The benchmarking tests outlined in this document are confined to evaluating the performance of SAV devices within a controlled laboratory environment, utilizing isolated networks.

The network topology employed for benchmarking must constitute an independent test setup. It is imperative that this setup remains disconnected from any devices that could potentially relay test traffic into an operational production network.

9. References

9.1. Normative References

- [RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, DOI 10.17487/RFC3704, March 2004, <<https://www.rfc-editor.org/rfc/rfc3704>>.
- [RFC8704] Sriram, K., Montgomery, D., and J. Haas, "Enhanced Feasible-Path Unicast Reverse Path Forwarding", BCP 84, RFC 8704, DOI 10.17487/RFC8704, February 2020, <<https://www.rfc-editor.org/rfc/rfc8704>>.
- [RFC2544] Bradner, S. and J. McQuaid, "Benchmarking Methodology for Network Interconnect Devices", RFC 2544, DOI 10.17487/RFC2544, March 1999, <<https://www.rfc-editor.org/rfc/rfc2544>>.
- [RFC4061] Manral, V., White, R., and A. Shaikh, "Benchmarking Basic OSPF Single Router Control Plane Convergence", RFC 4061, DOI 10.17487/RFC4061, April 2005, <<https://www.rfc-editor.org/rfc/rfc4061>>.
- [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/rfc/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/rfc/rfc8174>>.

9.2. Informative References

- [intra-domain-ps] "Source Address Validation in Intra-domain Networks Gap Analysis, Problem Statement, and Requirements", 2025, <<https://datatracker.ietf.org/doc/draft-ietf-savnet-intra-domain-problem-statement/>>.
- [inter-domain-ps] "Source Address Validation in Inter-domain Networks Gap Analysis, Problem Statement, and Requirements", 2025, <<https://datatracker.ietf.org/doc/draft-ietf-savnet-inter-domain-problem-statement/>>.
- [intra-domain-arch] "Intra-domain Source Address Validation (SAVNET) Architecture", 2025, <<https://datatracker.ietf.org/doc/draft-ietf-savnet-intra-domain-architecture/>>.

[inter-domain-arch]

"Inter-domain Source Address Validation (SAVNET)
Architecture", 2025, <[https://datatracker.ietf.org/doc/
draft-wu-savnet-inter-domain-architecture/](https://datatracker.ietf.org/doc/draft-wu-savnet-inter-domain-architecture/)>.

Acknowledgements

Many thanks to Aijun Wang, Nan Geng, Susan Hares, Giuseppe Fioccola, Minh-Ngoc Tran, Shengnan Yue, Changwang Lin, Yuanyuan Zhang, and Xueyan Song, for their valuable comments and reviews on this document. Apologies to any others whose names the authors may have missed mentioning.

Authors' Addresses

Li Chen
Zhongguancun Laboratory
Beijing
China
Email: lichen@zgclab.edu.cn

Dan Li
Tsinghua University
Beijing
China
Email: tolidan@tsinghua.edu.cn

Libin Liu
Zhongguancun Laboratory
Beijing
China
Email: liulb@zgclab.edu.cn

Lancheng Qin
Zhongguancun Laboratory
Beijing
China
Email: qinlc@zgclab.edu.cn