Journal of Computer and Communications, 2025, 13(6), 26-32
https://www.scirp.org/journal/jcc
ISSN Online: 2327-5227
ISSN Print: 2327-5219

IPsec Tunnel Recovery from Out-of-Sequence Traffic Drop Due to Peer IPsec Stateful Switchover

Arun Raj Kaprakattu

Nokia of America Corporation, Sunnyvale, CA, USA
Email: arunrajkaprakattu@gmail.com

How to cite this paper: Kaprakattu, A. (2025) IPsec Tunnel Recovery from Out-of-Sequence Traffic Drop Due to Peer IPsec Stateful Switchover. *Journal of Computer and Communications*, **13**, 26-32.
https://doi.org/10.4236/jcc.2025.136002

Received: May 21, 2025
Accepted: June 10, 2025
Published: June 13, 2025

Abstract

In an IPsec stateful high availability environment [1], synchronizing ESP sequence numbers between active and standby IPsec gateway devices is the most challenging, as these sequence numbers change with every ESP/AH packet [2]. Consider that the IPsec state synchronization occurs periodically. If a failover/switchover occurs, the difference between the last synchronized sequence number and the current sequence number on the active device exceeds the anti-replay window and then the standby device will not be aware of the last sent ESP/AH sequence number. Hence, traffic sent from the new active IPsec device uses outdated sequence numbers, which will get dropped by the remote IPsec peer if anti-replay mechanism is enabled [3]. The purpose of this document is to explain how a remote standalone IPsec peer can differentiate this from a replay attack and help a newly active IPsec peer recover from sending out-of-sequence ESP/AH traffic caused by a stateful switchover.

Keywords

Component, Formatting, Style, Styling, Insert (IPsec, Ike, Esp, Stateful Redundancy, Antireplay)

1. Introduction

IPsec [4] is a suite of protocols for securing IP connections between peers. It helps to secure the data sent over public networks. It works by encrypting IP packets, along with authenticating the source where the packets come from.

Let us consider a deployment scenario where one IPsec peer is running in standalone mode, while the remote IPsec peer is configured to operate in stateful redundancy. The sequence number [2] in IPsec is 32-bit field in the AH and ESP

headers and is used to: 1) ensure packets are received in order and 2) prevent attackers from replaying old packets (anti-replay detection). IPsec employs a windowing [5] mechanism to handle out-of-order packets.

This is demonstrated in Figure 1 below. Sequence numbers that fall on the left of the window are considered as past and dropped, while the ones within and on the right of the window are accepted.

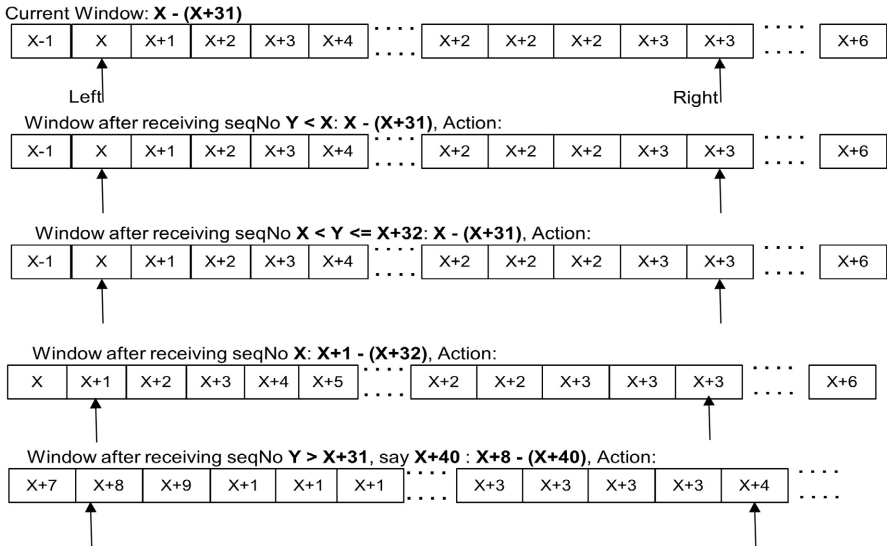

Figure 1. IPsec anti-replay window mechanism.

If an IPsec peer encounters an out-of-order sequence number falling outside the acceptable window, it will continue discarding the packets until the next rekey, either a Child SA rekey in IKEv2 [6] or a Phase2 rekey in IKEv1 [7], or the sequence number of the new stateful peer falls within the acceptable window of the standalone peer. Currently, no mechanism exists for the standalone IPsec peer to notify the redundant remote that it is receiving unacceptable out-of-sequence packets/sequence numbers.

1.1. Traffic Loss Estimation in Case of Failover

Consider a scenario where the synchronization interval (syncTime) between a stateful active and stateful standby peer is 10 seconds. If the active peer fails just before the next scheduled sync, the standby peer may not have the most recent state information. Assuming an approximate traffic rate of 140,000 packets per second (pps) per IPsec tunnel, and a total of 4,000 tunnels (tun count) on the chassis, the worst-case traffic loss can be estimated as:

TrafficLoss = syncTime × pps × tunCount = 10 × 140,000 × 4000 = 5.6 billion packets.

1.2. Information Synchronized between Active and Standby IPsec Stateful Chassis

IKE and IPsec Security Associations (SAs), mechanisms to handle replay counters

and message IDs to maintain session continuity, etc. are synchronized between active and standby chassis [8].

2. Scenarios That Can Cause Out-of-Sequence Traffic Drops in the Standalone IPsec Peer Due to Stateful Switchover of the Redundant Peer

2.1. High Availability of IPsec via Stateful Redundancy between Chassis

When two IPsec chassis are configured in a stateful redundancy, the IPsec state information is periodically synchronized from active to standby, as demonstrated in Figure 2.

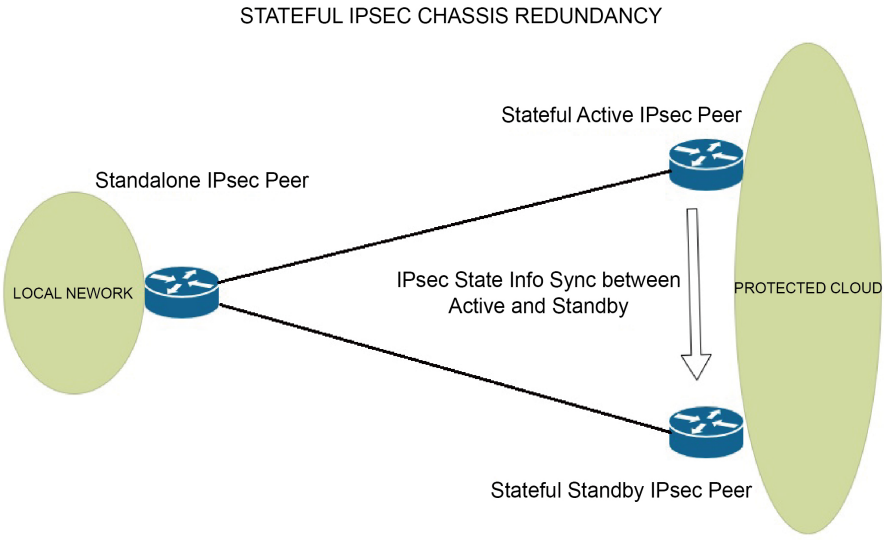

Figure 2. Stateful IPsec chassis-level redundancy.

However, this synchronization is not always perfectly accurate, especially the ESP/AH sequence number. There is often a calculated delay in the process, which can lead to inconsistency (information mismatch) between the active and standby during synchronizations. This can potentially cause packet loss or delayed traffic during failover. Considering there is a new sequence number for each ESP/AH packet, and if the peer is processing a high-throughput traffic, the packet loss can become significant. In a way, the amount of packet loss is directly proportional to both synchronization interval and the throughput.

2.2. Chassis with IPsec Stateful Line Card Level Redundancy

Compared to Section 2.1, here the synchronizing information is between line card levels within the same chassis, and communication must be done between line cards via the backplane. However, if the gateway is capable of high-throughput traffic, there will be a significant sequence number mismatch between each update.

3. Existing Solutions for Out-of-Sequence ESP/AH Drop after Stateful IPsec Switchover

3.1. Initiating a Phase2 Rekey by the New Stateful Redundant Peer Immediately after a Switchover

One of the existing solutions is that whenever a switchover occurs, the newly active device initiates a Phase 2 rekey immediately [8] (Here, Stateful Redundancy can be Line card level or node level), triggered by the redundant side. Once the new Phase2 Security Association (SA) is established, the sequence number will be in sync again.

Potential Issues with (3.1) rekey initiated by the redundant Gateway is that in most cases, the redundancy peer will be the aggregation gateway handles millions of IPsec tunnels. This can lead to significant processing and traffic issues, especially when all these tunnels attempt to rekey within a short period. The large volume of tunnels being rekeyed simultaneously can overwhelm the system, potentially causing delays, performance degradation, or even packet/tunnel loss during the rekeying process.

3.2. Disabling or Expanding Anti-Replay Window Protection

Disabling anti-replay or expanding anti-replay window [9] protection can help resolve the issue of out-of-sequence ESP/AH packet drops on the local side during a switchover in a redundant IPsec setup, specifically for data traffic. When anti-replay is disabled, the local peer does not perform sequence number checks, and it may accept out-of-sequence packets that would otherwise be discarded. This can help avoid packet drops caused by sequence number mismatches during failover or switchover events.

Potential Issues with (3.2) expanding anti-replay windows are that it makes the window too large and takes time to detect an attack. Disabling anti-replay window makes the tunnel susceptible to replay attacks as well as the receiver may accept out-of-order packets that go undetected. This can cause significant issues for applications like VoIP, video conferencing, etc.

4. The Proposed Solution

The proposed solution enables the standalone IPsec peer to initiate Phase2 rekey when it detects a switchover has occurred in the redundant peer.

For this, the standalone peer should first identify when a switchover has occurred in the stateful peer. It should be able to differentiate between an out-of-sequence drop caused by a real issue, like a network problem or a replay attack, and one caused by a stateful switchover on the peer.

This should be a feature to be enabled on the need basis on standalone peer side if IPsec redundancy is enabled remote peer. This feature consists of two components: 1) standalone peer detecting that a stateful switchover has occurred on the stateful peer, and 2) initiating a Phase 2/Child SA rekey from standalone peer based on IKE version used.

4.1. Detection of Stateful Switchover Events by a Standalone IPsec Peer at a Remote Endpoint

If only out-of-sequence ESP/AH packets are received on the inbound Security Association (SA) of the standalone IPsec peer for a defined interval or packet threshold—without any in-sequence packet observed during that period—this condition may indicate a remote Stateful Switchover event, it indicates a sequence number synchronization issue on the remote stateful peer. If it is a real anti-replay attack, there is a chance that some legitimate traffic also be received in between from the actual remote peer.

There is a scope to improve or add more ways to detect the stateful switchover by the standalone peer.

4.2. Inducing Phase 2/Child SA Rekey from a Standalone IPsec Peer upon Detection of Remote Redundancy Switchover

The proposed solution is to introduce functionality on the standalone side to initiate a Phase2 rekey (Phase2 rekey in the case of IKEv1, and Child SA rekey in the case of IKEv2) once a switchover is detected on the stateful peer.

4.3. Proposed Criteria for Detecting a Stateful Switchover and Initiating a Rekey from a Standalone Peer

- A Count-based Threshold for Incoming Out-of-sequence Packets: An induced rekey is initiated if the threshold number of out-of-sequence packets received on standalone peer, matching the incoming SA from the redundant peer exceeds a defined threshold. The threshold count will be reset if a single in-sequence packet is received in between to differentiate it from a real packet replay attack. This rekey is triggered based on the out-of-sequence threshold rather than the rekey timer. This threshold should be derived based on the synchronization interval between the stateful peers and the expected traffic rate within the IPsec tunnel. It must be user-configurable to accommodate varying deployment scenarios and traffic patterns.
- A Time-based Threshold for Incoming Out-of-sequence Packets: The time-based threshold is a pre-configured time on standalone peer where it only receives out-of-sequence traffic matching inbound SA from remote. This time should be less than configured Phase2/Child SA rekey time. The timer should be reset if a single in-sequence packet is received. This time should be less than the sync time between stateful peer and should be user configurable.
- Avoid Unwanted Induced Rekeys: To avoid unwanted induced rekeys from standalone peers there should be value that the maximum number of consecutive induced rekeys could be triggered within a certain time from a standalone peer. This threshold would help ensure that rekeying is initiated only when necessary, preventing excessive rekeying due to frequent IPsec switchover on Remote due to a fault or a replay attack compromising the anti-replay pattern during remote switchover.

After reaching maximum induced rekey count within a certain time, there should be a quiet period during which the standalone device does not initiate any further induced rekey based on the aforementioned condition. During this period, standalone peers should ignore any pattern that matches the stateful switchover if identified. Once the quiet period expires, this timer should be reset, allowing the device to resume its operation based on the reception of out-of-sequence packets. This ensures that the system can recover from transient issues while maintaining security by filtering out potential replay attacks as well as a fault that causes the remote to do frequent switchovers.

- Never induce a rekey for out-of-sequence recovery from a standalone peer side if out-of-sequence packets are arriving in the older inbound SA, while a newer inbound SA has just been installed as part of a recently completed rekey. This new security association may have been created either through rekeying initiated by the redundant gateway or because of a time-based rekeying event.

5. Comparison of the Induce Rekey Solution vs. Rekey Initiated by Redundant Peer

5.1. Newly Active Redundant Peer Initiating Phase2/Child SA Rekey for All Tunnel after Switchover to Avoid Out-of-Sequence Traffic

In this approach, all tunnels in the newly active stateful peer need to initiate a Phase2/Child SA rekey [8], which is unnecessary and creates a processing spike in the gateway router, irrespective of whether the remote is facing out-of-sequence drops or not. Even if the redundant gateway tries to identify which peers are experiencing out-of-sequence drops, it's impossible. There are several other factors that make a tunnel unaffected by out-of-sequence traffic issues, such as the timing of switchover, comparatively lower traffic rates, etc. An induced rekey for those tunnels is unnecessary.

Whereas the proposed solution requires only the peers that receive out-of-sequence IPsec packets to initiate the rekey. Tunnels that fall within the sequence number windowing mechanism do not need to be rekeyed.

5.2. Disabling Anti-Replay

Disabling anti-replay can affect real-time traffic, such as voice and video, which may experience increased jitter as well as susceptible to anti-replay attacks. With the proposed solution in Section 3, the IPsec devices can still enable anti-replay protection, effectively avoiding replay attacks while addressing the issue of out-of-sequence packets issue due to switchover.

6. Conclusion

This is an alternative method compared to what is described in the Strong Swan High Availability documentation [10]. This approach offers advantages compared to rekeying all active tunnels in a stateful system, as it allows for selective rekeying

based on out-of-window drops. More research is needed in Section 4.1 regarding the identification of a switchover event in the redundant peer.

Conflicts of Interest

The author declares no conflicts of interest regarding the publication of this paper.

References

[1] Kuboniwa, A., Tamura, T., Huruta, Y., Wada, Y., Satou, H. and Motono, T. (2010) IPsec-GW Redundancy Method with High Reliability. 8*th Asia-Pacific Symposium on Information and Telecommunication Technologies*, Kuching, 15-18 June 2010, 1-5. https://ieeexplore.ieee.org/abstract/document/5532017

[2] Kent, S. (2005) IP Encapsulating Security Payload (ESP). https://www.rfc-editor.org/rfc/rfc4303

[3] Gouda, M., Huang, C.-T. and Li, E. (2000) Anti-Replay Window Protocols for Secure IP. *Proceedings Ninth International Conference on Computer Communications and Networks*, Las Vegas, 16-18 October 2000, 310-315. https://doi.org/10.1109/icccn.2000.885507

[4] Dunbar, N. (2001) IPsec Networking Standards—An Overview. *Information Security Technical Report*, **6**, 35-48. https://doi.org/10.1016/s1363-4127(01)00106-6

[5] Zhao, F. and Wu, S.F. (2003) Analysis and Improvement on IPSec Anti-Replay Window Protocol. *Proceedings of* 12*th International Conference on Computer Communications and Networks*, Dallas, 22-22 October 2003, 553-558. https://doi.org/10.1109/icccn.2003.1284223

[6] Kaufman, C. (2005) Internet Key Exchange (IKEv2) Protocol. https://www.rfc-editor.org/rfc/rfc4306

[7] Harkins, D. and Carrel, D. (1998) The Internet Key Exchange (IKE). https://www.rfc-editor.org/info/rfc2409

[8] Kalyani, G., Singh, R., Nir, Y., Sheffer, Y. and Zhang, D. (2011) Protocol Support for High Availability of IKEv2/IPsec. https://www.rfc-editor.org/rfc/rfc6311

[9] Cisco (2008) IPsec Anti-Replay Window Expanding and Disabling. Cisco. https://www.cisco.com/c/en/us/td/docs/routers/asr920/configuration/guide/sec_vpn/sec-ipsec-xe-3s-book-920/configuring_ipsec_anti_replay_window_expanding_and_disabling.pdf

[10] Ullah, S., Choi, J. and Oh, H. (2020) IPsec for High Speed Network Links: Performance Analysis and Enhancements. *Future Generation Computer Systems*, **107**, 112-125. https://doi.org/10.1016/j.future.2020.01.049