Protocol Support for High Availability of IKEv2/Ipsec

draft-ietf-ipsecme-ipsecha-protocol-02

Dacheng Zhang

IPsecME WG

IETF-79, Beijing

v00

Motivation Scenario

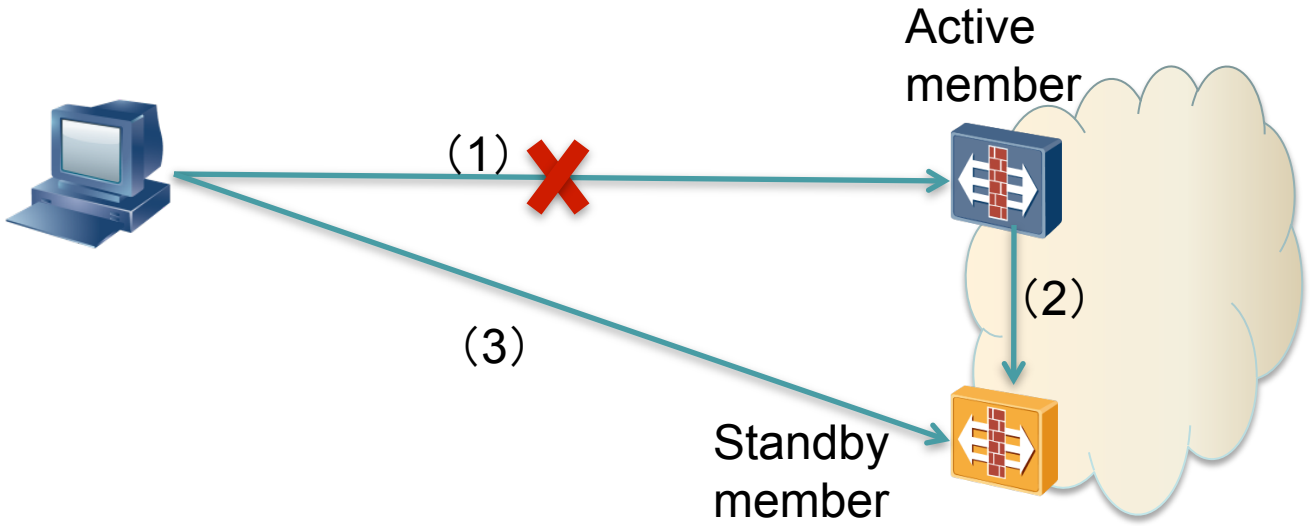

1. A user establishes an IKEv2/IPsec session with the active member of a hot standby cluster
2. The active member syncs IKEv2/IPsec SA states to the standby member
3. A “failover” event occurs in cluster; the standby member thus takes over the failed one and becomes the active member

State Synchronizing Issue

- In a failover event, the state information held by standby members may be stale:
  - Before an active member fails, it may be not able to send the latest state information to the corresponding standby members
  - For the purpose of efficiency, state information may only be synchronized in a periodic way
  - New packets may have been sent out by users during the period of the failover

Scope of the Draft

- This draft attempts to address the issues mentioned in [RFC6027]:
  - IKE Counters
  - Outbound SA Counters
  - Inbound SA Counters
  - Missing Synchronization Messages
- Only “tight” IPsec clusters are considered
- Both the synchronization issues with IKEv2 Message ID counters and IPsec replay counters are considered

Solution (1)

The new active member needs to synchronize KEv2 Message ID counters, and IPsec replay counters with the user

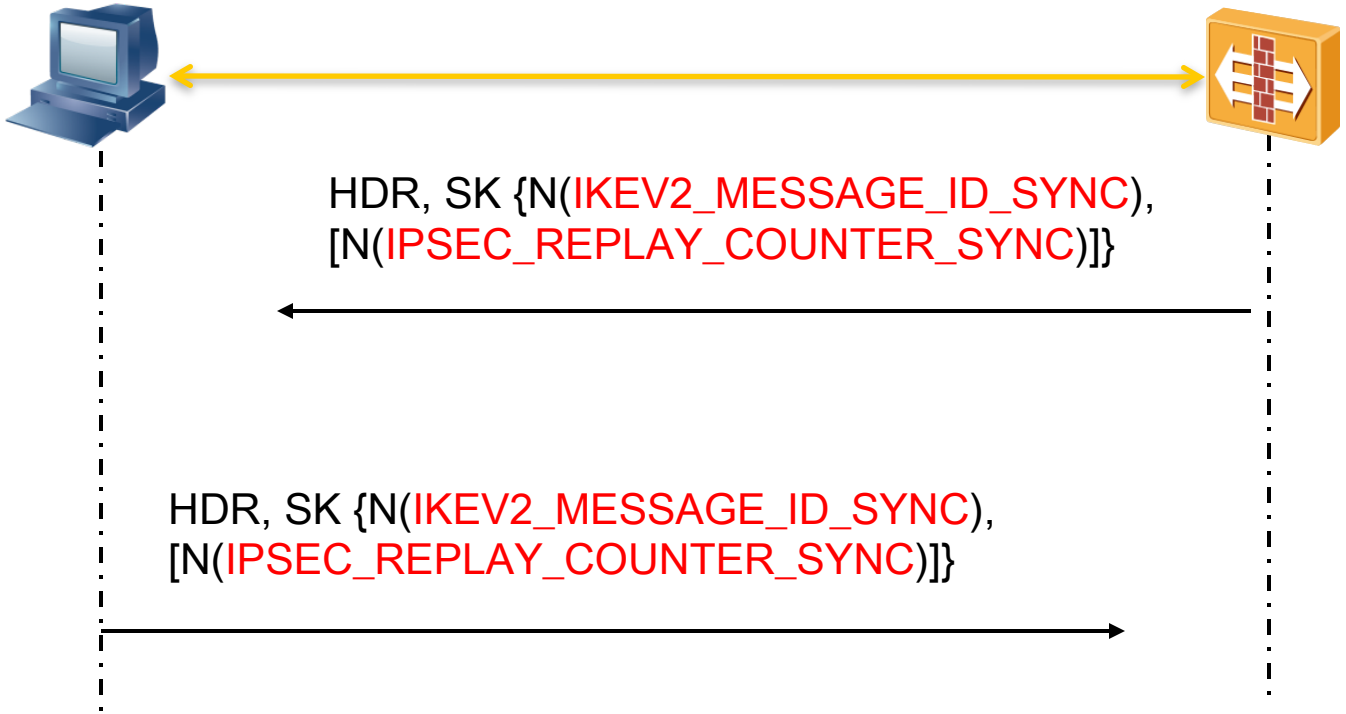

IKEV2_MESSAGE_ID_SYNC

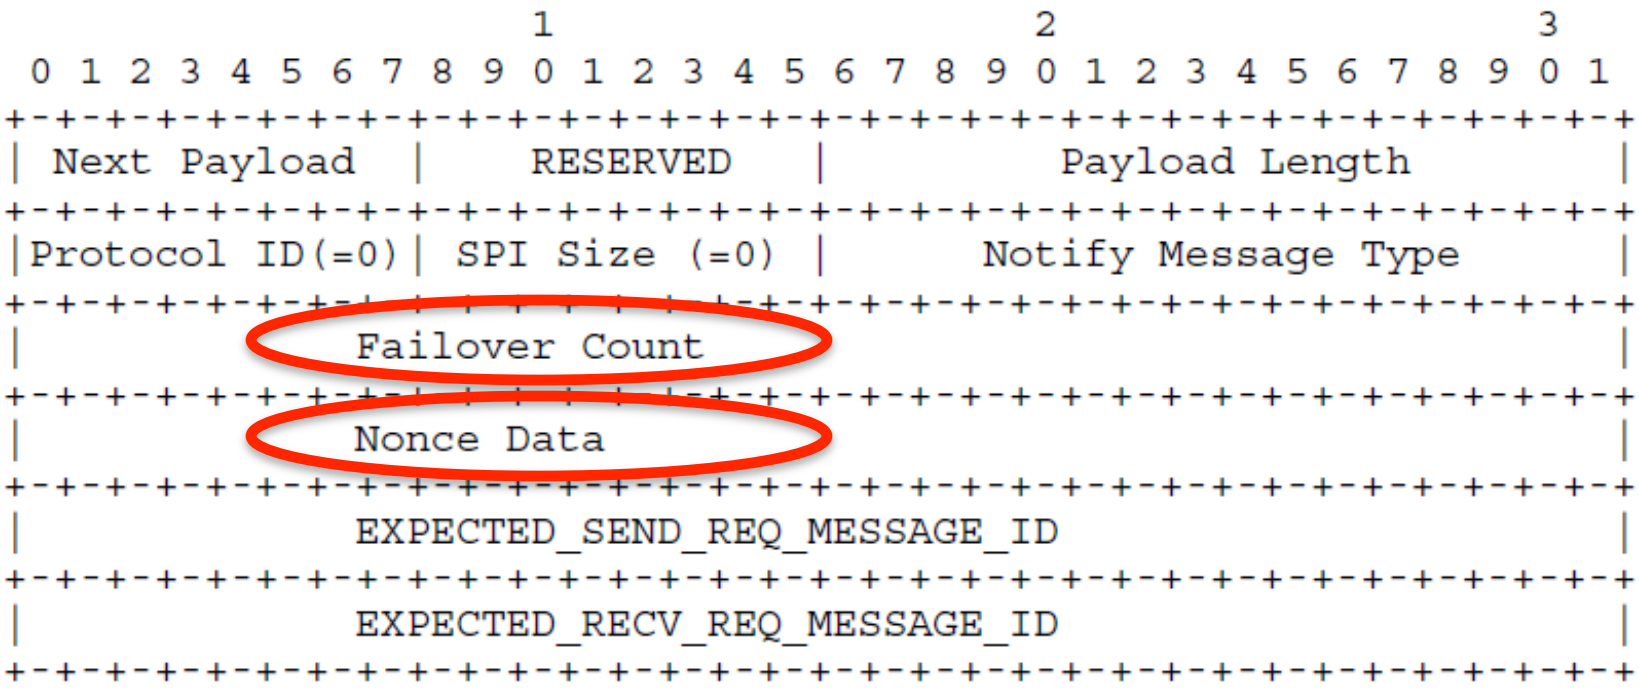

State Synchronization in Simultaneous Failover Scenarios

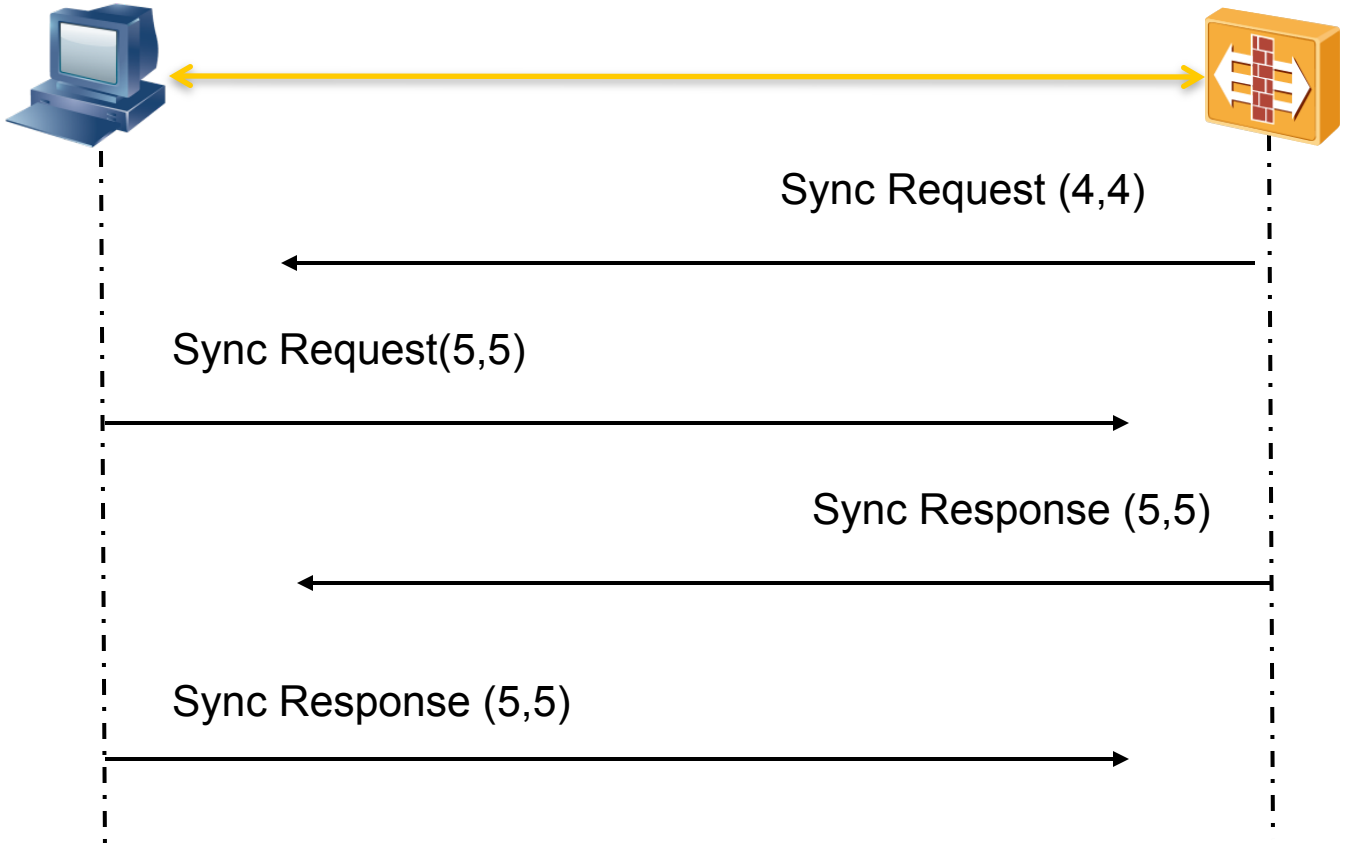

IPSEC_REPLAY_COUNTER_SYNC

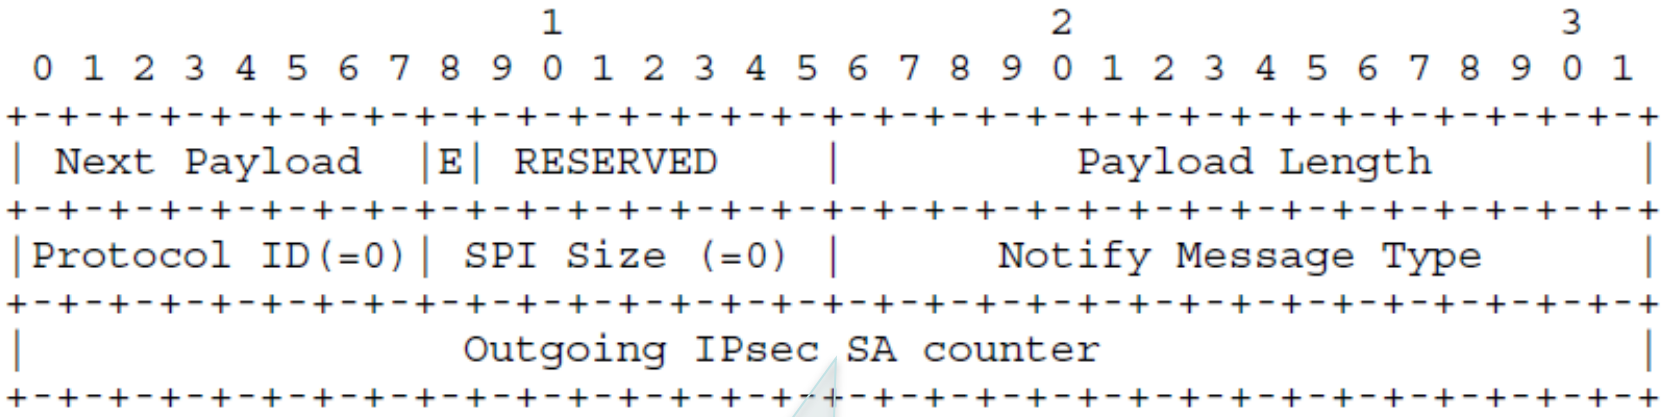

Used to transport count delta value actually

Solution (2)

- The user needs to negotiate the ability to sync SA counters in their original IKE_AUTH exchange

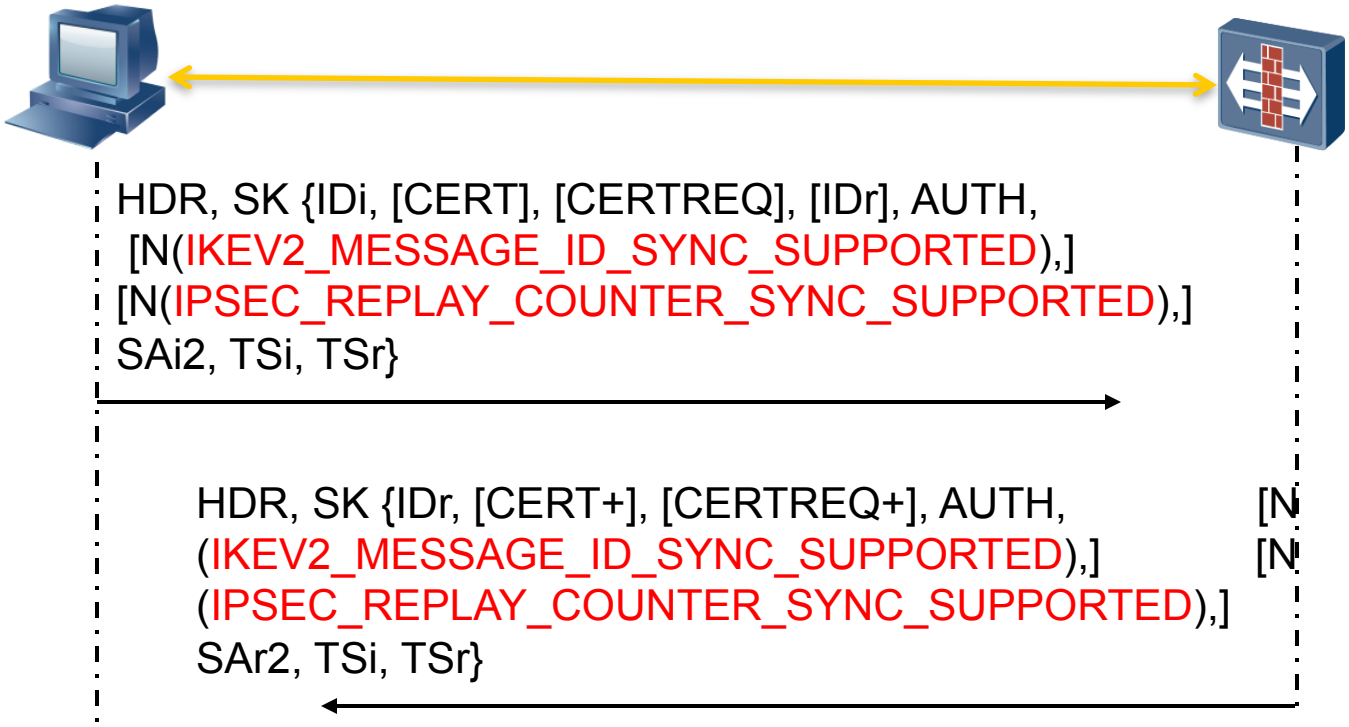

Sync Support Notifications

- IKEV2_MESSAGE_ID_SYNC_SUPPORTED

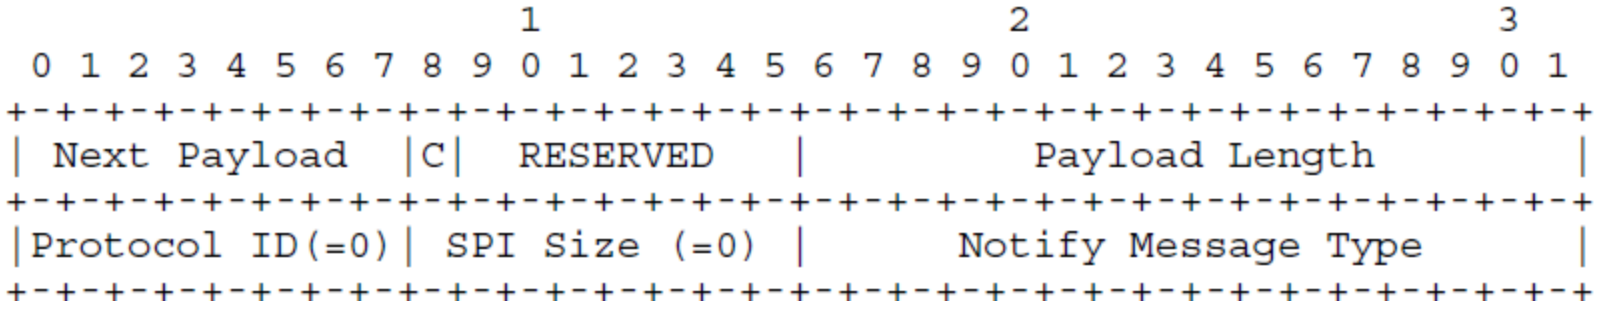

- IPSEC_REPLAY_COUNTER_SYNC_SUPPORTED

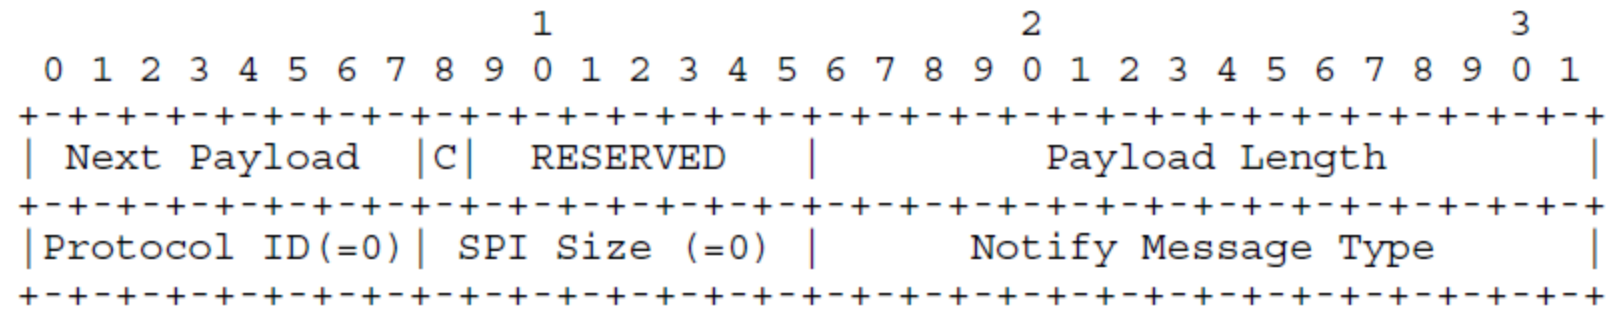

Resistance to Replay Attacks

- Replay of Message ID synchronization requests:
  - The receiver of the synchronization request should verify the received Failover Count and maintain its own copy of it. If a peer receives a synchronization request with an already observed Failover Count, it can safely discard the request
- Replay of Message ID synchronization responses:
  - This is countered by sending the nonce data along with the synchronization payload.

Major Differences from Previous Versions

- Introduce failover counts to resist replay attacks
- The sync of IPsec SA replay counter is optimized to have just one global bumped-up outgoing IPsec SA counter delta of ALL Child SAs under an IKEv2 SA
- The proposed mechanism is now able to sync either IKEv2 message ID, IPsec replay counter, or both to cater different types of implementations.
- Demonstrate how the mechanism works in the multiple and simultaneous failover scenarios

Issues (1)

- Multiple failover: which is the situation where, in a cluster with three or more members, failover happens in rapid succession. It is our goal that the implementation should be able to handle this situation, i.e. to handle the new failover event even if it is still processing the old failover

Issues (2)

- How to synchronize the failover counter amongst different cluster members
  - Multiple Failure Scenarios
- How to transport the latest the failover count value to a user when it initially access an active member

Next Step

- Before we are ready for WG Last Call, follows should be done
  - Collecting comments in the list
  - Solve the issues known so far
  - ……

Comments?