

Aggregated Option for SYN Option Space Extension

draft-nishida-tcpm-agg-syn-ext-03

Yoshifumi Nishida
nsd+ietf@gmail.com

Purpose

- Provide a solution for limited option space in SYN segments
 - All options to be used in a connection should be appeared in SYN
 - Otherwise, the features should not be used

Previous Approach for SYN Option Space Extension

- Extending header format for SYN segments
 - e.g. Override DataOffset field in TCP header
- Sending multiple SYN or SYN-like segments
 - Use multiple segments to carry more options
- Combination of above

Our Approach to Save Option Space in SYN

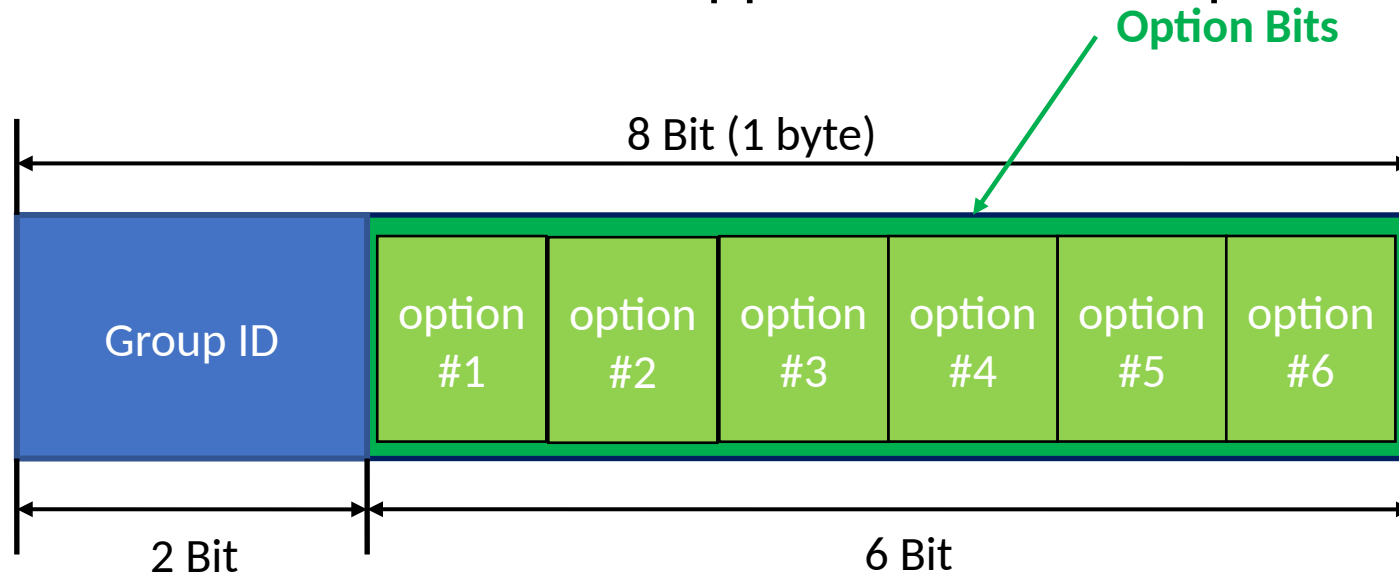
- Define “Aggregated Option”
 - A new TCP option used in only SYN segments
- Don't modify header format
- Don't use additional segments
- Don't change TCP's semantics

Concept for Aggregated Option

- Some options in SYN are just used for “indication of the feature”
 - e.g. SACK-Permit (RFC2018), Fast Open Cookie Request (RFC7413), EDO (draft-ietf-tcpm-tcp-edo), TARR (draft-tcpm-ack-rate-request)
- This is 1 bit information
 - However, each option consumes 2 or 4 bytes or more option space
 - Experimental Options (RFC6994) will have at least 4 bytes length
- Aggregating them into a single option can create more option space

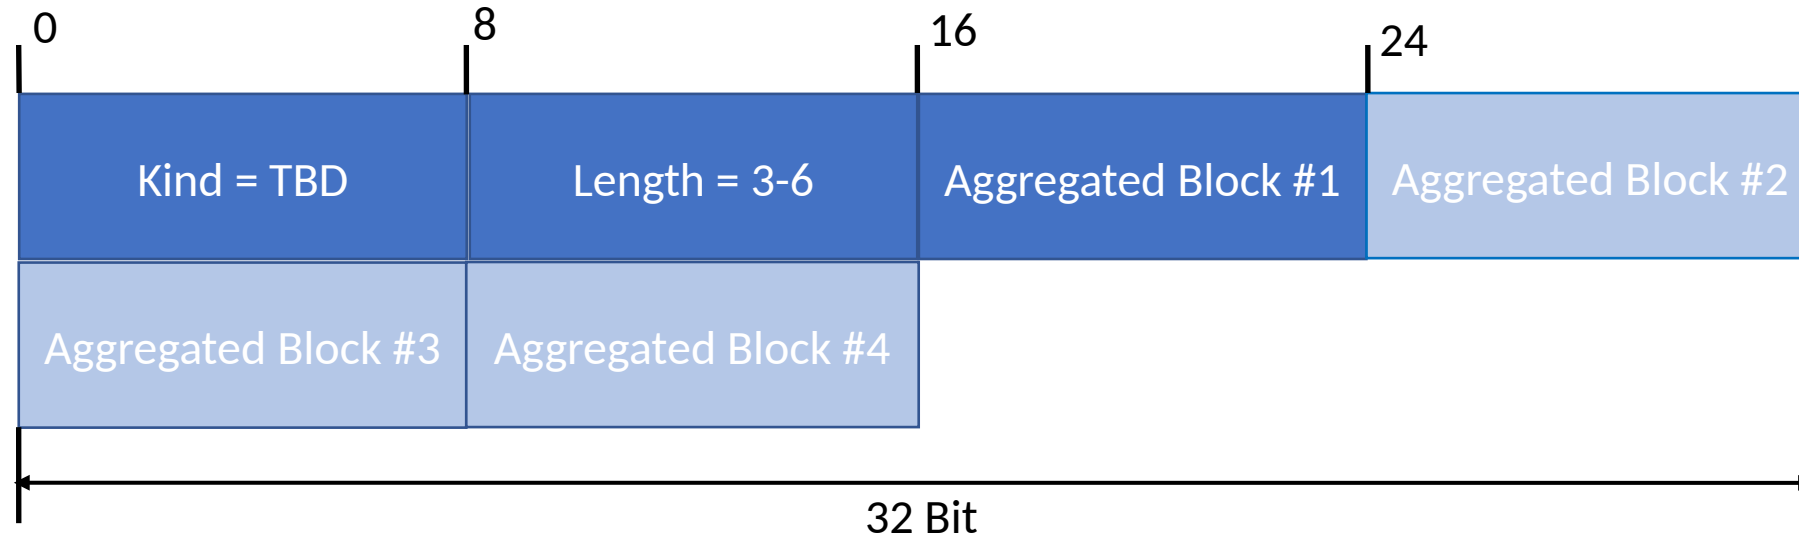
Aggregated Block Format

- ‘Aggregated block’ is used to compress multiple options
 - **Group ID** has 2 bits, **Option Bits** has 6 bits
 - Each bit in **Option Bits** represents one TCP option kind
 - E.g. 3rd bit of Group ID 0 indicates X option, 5th bit of Group ID 1 indicates Y option
 - Mapping between Option Bit and TCP Option kinds will be maintained by IANA
- As a results, this format can support $4 \times 6 = 24$ options

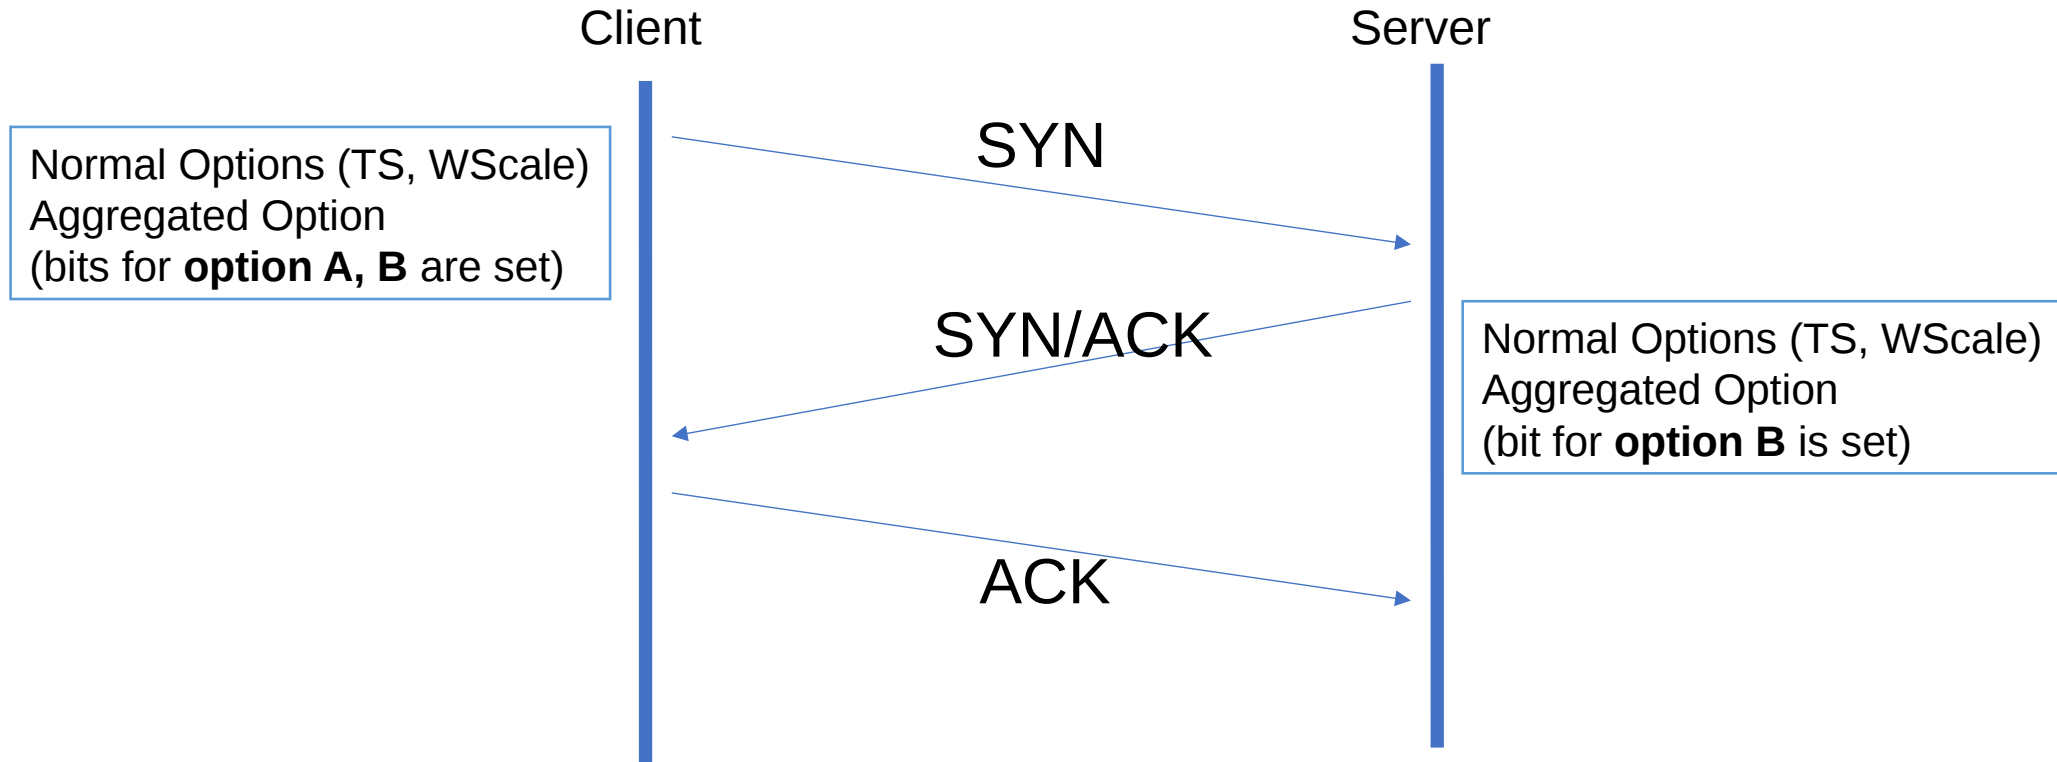


Aggregated Option Format

- **Aggregated Option** can contain 1-4 Aggregated blocks
 - Each block has 1 byte length



Aggregated Option Example Usage



- In this example, Option B will be used in the connection through Aggregated Option

Deployment Strategies

- Deploy Aggregated Option with a newly defined option
 - Ask implementors to support aggregated option when they implement a new feature
- Send original option and Aggregated option in SYN to new destination
 - If Aggregated option is in SYN/ACK, cache the info for future connections

Updates from -00 version

- It contains 'Aggregated Option' only
 - There will be use cases for some new features (e.g. TARR)
- Removed 'Delayed Option Negotiation' Scheme for simplicity
 - We currently do not have a concrete use case for Delayed Option Negotiation

Conclusion

- Aggregate multiple TCP options to save option space in SYN
- We believe this is a practical approach for SYN option space issue
 - If compress approach is not preferable, need to consider approach like extending header format, sending multiple segments
 - If needed, we can try both approach
 - At least, our approach won't create conflicts as it doesn't change TCP semantics