

sidrops  
Internet-Draft  
Intended status: Standards Track  
Expires: 2 October 2026

K. Xu  
X. Wang  
Z. Liu  
Q. Li  
J. Wu  
Tsinghua University  
Y. Guo  
Zhongguancun Laboratory  
31 March 2026

FC-BGP Protocol Specification  
draft-wang-sidrops-fcbgp-protocol-05

## Abstract

This document defines an extension, Forwarding Commitment BGP (FC-BGP), to the Border Gateway Protocol (BGP). FC-BGP provides security for the path of Autonomous Systems (ASs) through which a BGP UPDATE message passes. Forwarding Commitment (FC) is a cryptographically signed segment to certify an AS's routing intent on its directly connected hops. Based on FC, FC-BGP aims to build a secure inter-domain system that can simultaneously authenticate the AS\_PATH attribute in the BGP UPDATE message and alleviate route leaks in the BGP routing system. The extension is backward compatible, which means a router that supports the extension can interoperate with a router that doesn't support the extension.

## About This Document

This note is to be removed before publishing as an RFC.

The latest revision of this draft can be found at <https://FCBGP.github.io/fcbgp-protocol/draft-wang-sidrops-fcbgp-protocol.html>. Status information for this document may be found at <https://datatracker.ietf.org/doc/draft-wang-sidrops-fcbgp-protocol/>.

Source for this draft and an issue tracker can be found at <https://github.com/FCBGP/fcbgp-protocol>.

## 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 2 October 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. Requirements Language . . . . .	5
1.2. Definitions, and Acronyms . . . . .	5
2. FC-BGP Negotiation . . . . .	6
3. FC Path Attribute . . . . .	6
4. FC-BGP UPDATE Messages . . . . .	10
4.1. Generation . . . . .	10
4.2. Propagation . . . . .	13
4.3. Processing Instructions for AS Confederation Members . .	14
4.4. Processing Instructions for BGP Route Leak Prevention . .	15
5. Processing a Received FC-BGP UPDATE Message . . . . .	16
5.1. Overview . . . . .	16
5.2. Validation . . . . .	16
5.2.1. Validation Algorithm . . . . .	18
6. Implementations, Operations, and Management Considerations .	21
6.1. Algorithms and Extensibility . . . . .	21
6.2. Speedup and Early Termination of Signature Verification . . . . .	22
6.3. Deferring Validation . . . . .	23
6.4. BGP Route Selection . . . . .	23

6.5.	Non-deterministic Signature Algorithms . . . . .	24
6.6.	Private AS Numbers . . . . .	25
6.7.	Robustness Considerations for Accessing RPKI Data . . . . .	25
6.8.	Graceful Restart . . . . .	25
6.9.	Robustness of Secret Random Number in ECDSA . . . . .	25
6.10.	Incremental/Partial Deployment Considerations . . . . .	25
6.11.	Co-existence with BGPsec . . . . .	26
7.	Security Considerations . . . . .	27
7.1.	Security Guarantees . . . . .	27
7.2.	Mitigation of Denial-of-Service Attacks . . . . .	27
7.3.	Route Server . . . . .	28
7.4.	Additional Security Considerations . . . . .	28
7.4.1.	Three AS Numbers . . . . .	28
7.4.2.	MISC . . . . .	29
8.	IANA Considerations . . . . .	29
9.	References . . . . .	29
9.1.	Normative References . . . . .	29
9.2.	Informative References . . . . .	31
Appendix A.	Attachment . . . . .	33
A.1.	Comparison to Other Technologies . . . . .	33
A.1.1.	BGPsec . . . . .	33
A.1.2.	ASPA . . . . .	34
A.1.3.	Only to Customer (OTC) Attribute . . . . .	34
A.2.	Implementation Status . . . . .	35
A.3.	An Example . . . . .	35
	Acknowledgments . . . . .	37
	Authors' Addresses . . . . .	37

## 1. Introduction

In a post-ROV (Route Origin Validation) era, route hijacks are largely mitigated but not eliminated, which fundamentally shifts the security focus up the stack. The problem space becomes less about "who owns the prefix" and more about "how routes propagate and whether that propagation is trustworthy and policy-compliant".

The FC-BGP mechanism described in this document aims to ensure that advertised routes in BGP [RFC4271] are authentic and alleviate the BGP route leaks.

FC-BGP accomplishes this by introducing a new optional, transitive, and extended-length path attribute called FC (Forwarding Commitment) to the BGP UPDATE message. This attribute can be used by an FC-BGP-compliant BGP speaker (hereafter referred to as an FC-BGP speaker) to generate, propagate, and validate BGP UPDATE messages to enhance security. In other words, when the BGP UPDATE message travels through an FC-BGP-enabled autonomous system (AS), it adds a new FC based on the AS order in AS\_PATH. Subsequent ASs can then utilize

the list of FCs in the BGP UPDATE message to ensure that the advertised path is consistent with the AS\_PATH attribute. And as a complementary of [RFC9234], it can also alleviate the BGP route leaks.

BGPsec is a path-level authentication approach described in [RFC8205]. It replaces the AS\_PATH attribute, which is used to record the sequence of ASs that a BGP update has traversed, with the non-transitive BGPsec\_Path attribute. However, when a peer does not support BGPsec, the BGPsec\_Path attribute will be downgraded to the standard AS\_PATH attribute, losing the security benefits BGPsec provides. In contrast, FC-BGP (Forwarding Commitment BGP) preserves the AS\_PATH attribute and introduces an additional list of signed messages called Forwarding Commitments. Each Forwarding Commitment (FC) is a publicly verifiable code certifying the correctness of a three-hop pathlet. FC-BGP builds its path authentication based on these FCs.

FC-BGP and BGPsec offer different levels of security benefits in the case of partial deployment, even though they achieve the same security benefits when fully deployed. BGPsec tightly couples path authentication with the BGP path construction process, requiring each AS to iteratively verify the signatures of each prior hop before extending the authentication chain. Consequently, a single legacy AS that does not support BGPsec can break the authentication chain, preventing subsequent BGPsec-aware ASs from reviving the authentication process. As a result, in partial deployment scenarios, BGPsec is often downgraded to the legacy BGP protocol, losing its security benefits.

In contrast to BGPsec, FC-BGP treats partial deployability as a first-class citizen. It adopts a pathlet-driven authentication paradigm, in which the authenticity of an AS path can be incrementally built based on authenticated pathlets. This design ensures that downstream FC-BGP-aware ASs can use the authenticated pathlets provided by upstream upgraded ASs, even if the full AS path traverses legacy ASs that do not support FC-BGP. By allowing the authentication of sub-paths, FC-BGP enables incremental deployment and provides security benefits to the FC-BGP-aware ASs, regardless of the deployment status of other ASs on the path Appendix A.1.1.1.

Similar to BGPsec, FC-BGP relies on RPKI to perform route origin validation [RFC6483]. Additionally, any FC-BGP speaker that wishes to process the FC path attribute along with BGP UPDATE messages MUST obtain a router certificate and store it in the RPKI repository. This certificate is associated with its AS number. The router key generation here follows [RFC8208] and [RFC8635].

It is NOT RECOMMENDED that both BGPsec and FC-BGP be enabled simultaneously in a BGP network. However, if a BGP update message contains both BGPsec and FC-BGP features, the BGP speaker should process the message properly. In such cases, the BGP speaker should prioritize BGPsec over FC-BGP. This means that if a BGP update message includes the BGPsec\_PATH attribute, a BGP speaker that supports both BGPsec and FC-BGP should use the Secure\_Path instead of the AS\_PATH to generate or verify the FC segments. This prioritization ensures that the presence of FC-BGP does not compromise the security benefits of BGPsec in the same update message. More discussion is at Section 6.11.

### 1.1. 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.

### 1.2. Definitions, and Acronyms

The following terms are used with a specific meaning:

**BGP neighbor:**

Also just 'neighbor'. Two BGP speakers that communicate using the BGP protocol are neighbors. It can be divided into iBGP neighbors and eBGP neighbors.

**BGP speaker:**

A device, usually a router, exchanges routes with other BGP speakers using the BGP protocol.

**BGP UPDATE:**

The message is generated with several path attributes to advertise routes.

**iBGP:**

iBGP neighbor, internal BGP neighbor, or internal neighbor. Internal neighbors are in the same AS.

**eBGP:**

eBGP neighbor, external BGP neighbor, or external neighbor. External neighbors are in different ASs.

**Router:**

In this document, the router always refers to a BGP speaker.

In addition to the list above, the following terms are used in this document:

**FC:**

Forwarding Commitment, i.e., FC segment. It contains several fields and a digital signature to protect the current path.

**FCList:**

An ordered list of FC segments to protect the whole AS-Path in the BGP UPDATE message. The order of FCs follows the order of AS numbers in the AS-Path. All FC-BGP-enabled BGP speakers SHOULD add their FCs to the BGP UPDATE message.

**FC path attribute:**

The optional, transitive, extended length path attribute is defined in this document to obtain BGP security.

**FC-BGP UPDATE:**

A BGP UPDATE message carries the FC path attribute.

**FC-BGP speaker:**

A BGP speaker that enables the FC-BGP feature. It can generate, propagate, and validate FC-BGP UPDATE messages.

## 2. FC-BGP Negotiation

FC-BGP does not need to negotiate with neighbors since it is considered a transitive path attribute within the BGP UPDATE message. BGP speakers that do not recognize the FC path attribute or do not support FC-BGP SHOULD still transmit the FC path attribute to their neighbors. As a result, there is no need to establish a new BGP capability as defined in [RFC5492].

However, if one AS has uploaded its keys to RPKI, it would be deemed to support FC-BGP. The reason and process are defined in Section 5.2.

## 3. FC Path Attribute

Unlike BGPsec, FC-BGP does not modify the AS\_PATH. Instead, FC is enclosed in a BGP UPDATE message as an optional, transitive, and extended-length path attribute. This document registers a new attribute type code for this attribute: TBD, see Section 8 for more information.

The FC path attribute includes the digital signatures that protect the pathlet information. We refer to those update messages that contain the FC path attribute as "FC-BGP UPDATE messages". Although

FC-BGP would not modify the AS\_PATH path attribute, it is REQUIRED to never use the AS\_SET or AS\_CONFED\_SET in FC-BGP according to [RFC6472] and [Deprecation-AS\_SET-AS\_CONFED\_SET].

The format of the FC path attribute is shown in Figure 1 and Figure 2. Figure 1 shows the format of FC path attribute and Figure 2 shows the FC segment format.

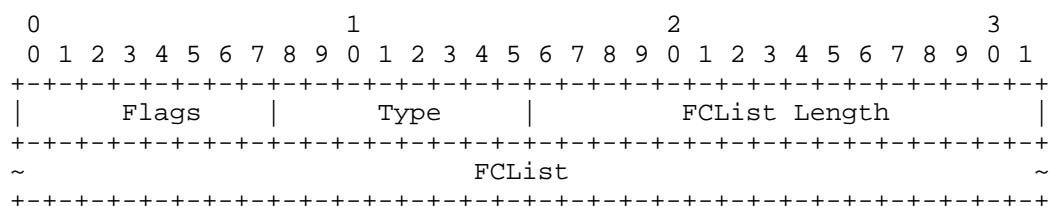


Figure 1: Format of FC path attribute.

FC path attribute includes the following parts:

Flags (1 octet):

The current value is 0b11010000, representing the FC path attribute as optional, transitive, partial, and extended-length.

Type (1 octet):

The current value is TBD. Refer to Section 8 for more information.

FCList Length (2 octets):

The value is the total length of the FCList in octets.

FCList (variable length):

The value is a sequence of FC segments, in order. The definition of the FC segment format is shown in Figure 2. It does not conflict with the AS\_PATH attribute. That means the FC-BGP path attribute and AS\_PATH attribute can coexist in the same FC-BGP UPDATE message.

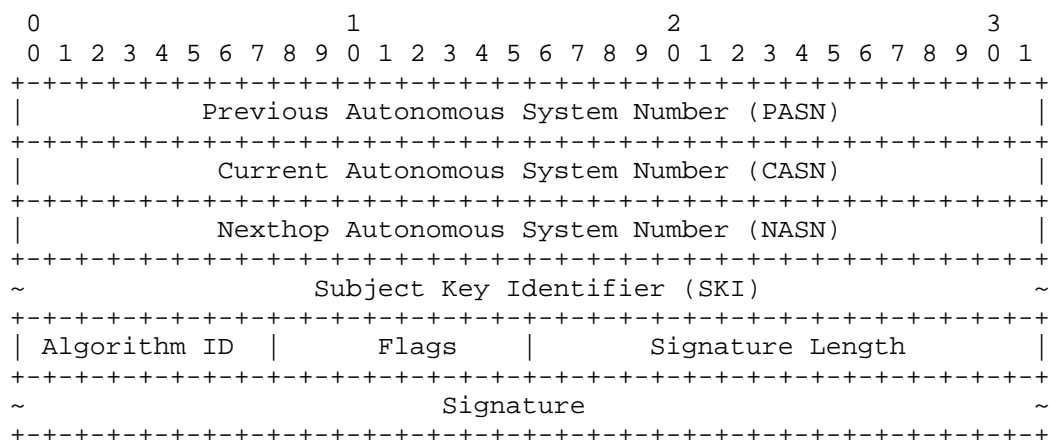


Figure 2: Format of FC segment.

In FC-BGP, all ASs MUST use 4-byte AS numbers in FC segments. Existing 2-byte AS numbers are converted into 4-byte AS numbers by setting the two high-order octets of the 4-octet field to 0 [RFC6793].

FC segment includes the following parts. (See Section 4 for more details on populating these fields.)

Previous Autonomous System Number (PASN, 4 octets):

The PASN is the AS number of the previous hop AS from which the FC-BGP speaker receives the FC-BGP UPDATE message. If the current AS has no previous AS hop, it MUST be filled with 0. It will be discussed more at Section 7.4.1.

Current Autonomous System Number (CASN, 4 octets):

The CASN is the AS number of the FC-BGP speaker that added this FC segment to the FC path attribute.

Nexthop Autonomous System Number (NASN, 4 octets):

The NASN is the AS number of the next hop AS to which the FC-BGP speaker will send the BGP UPDATE message.

Subject Key Identifier (SKI, 20 octets):

The SKI in the RPKI router certificate is a unique identifier for the public key used for signature verification. If the SKI length exceeds 20 octets, it should retrieve the leftmost 20 octets.

Algorithm ID (1 octet):

The current assigned value is 1, indicating that SHA256 is used to hash the content to be signed, and ECDSA is used for signing. It follows the algorithm suite defined in [RFC8208] and its updates. Each FC segment has an Algorithm ID, so there is no need to worry about sudden changes in its algorithm suite. The key in FC-BGP uses the BGPsec Router Key, so its generation and management follow [RFC8635].

Flags (1 octet):

Several flag bits. The leftmost bit of the Flags field in Figure 2 is the Confed\_Segment flag (Flags-CS). The Flags-CS flag is set to 1 to indicate that the FC-BGP speaker that constructed this FC segment is sending the UPDATE message to a peer AS within the same AS confederation [RFC5065]. (That is, a sequence of consecutive Confed\_Segment flags is set in an FC-BGP UPDATE message whenever, in a non-FC-BGP UPDATE message, an AS\_PATH segment of type AS\_CONFED\_SEQUENCE occurs.) In all other cases, the Flags-CS flag is set to 0. The second leftmost bit (i.e., the second highest) of the Flags field in Figure 2 is the Route\_Server flag (Flags-RS). The Flags-RS flag is set to 1 to indicate that a route server adds this FC segment, but the AS number will never appear in the AS\_PATH attribute. If the AS number of a router server is inserted into AS\_PATH, this Flags-RS flag MUST be set to 0. The third leftmost bit (i.e., the third highest bit) of the Flags field in Figure 2 is the AS\_Path\_Prepending flag (Flags-ASPP). This Flags-ASPP flag is set to 1 when AS path prepending is used; otherwise, this flag MUST be set to 0. The fourth and fifth leftmost bits (i.e., the fourth highest and fifth highest bits) of the Flags field in Figure 2 are the Provider\_to\_Customer flag (Flags-P2C) and Peer\_to\_Peer flag (Flags-P2P) separately. The Flags-P2C flag is set to 1 to indicate that the FC segment's issuer AS sends routes to its customer, and the Flags-P2P flag is set to 1 to indicate that the FC segment's issuer AS sends routes to its peer. If the Flags-P2C or Flags-P2P flags are set, the next route propagation will only be permitted to the following customers. The remaining bits of the Flags field are unassigned. They MUST be set to 0 by the sender and ignored by the receiver.

Signature Length (2 octets):

It only contains the length of the Signature field in octets, not including other fields.

Signature (variable length):

The signature content and order are `Signature=ECDSA(SHA256(PASN, CASN, NASN, SKI, Algorithm ID, Flags, Signature Length, Prefix, Prefix Length))`, where the Signature Length is kept as 0, the Prefix is the IP address prefix which is encapsulated in the BGP UPDATE, i.e., NLRI, and only one prefix is used each time. When hashing and signing, the full IP address and IP prefix length are used, i.e., IPv4 uses 4 octets, and IPv6 uses 16 octets.

#### 4. FC-BGP UPDATE Messages

##### 4.1. Generation

This part defines the generation of the FC path attribute and the FC segment. An FC-BGP speaker SHOULD generate a new FC segment or even a new FC path attribute when it propagates a route to its external neighbors. For internal neighbors, the FC path attribute in the BGP UPDATE message remains unchanged.

The FC-BGP speaker follows a specific process to create the FC path attribute for the ongoing UPDATE message. Firstly, it generates a new FC Segment containing the necessary information for the FC path. This new FC Segment is then added to the FCList of the outer format defined in Figure 1. It is important to note that if the FC-BGP speaker is not the origin AS and there is already an existing FC path attribute of the UPDATE message, it MUST prepend its new FC segment to the FCList of the existing FC path attribute, like the insertion process of ASN in the AS\_PATH attribute. This allows the FC-BGP speaker to contribute to its own FC segment while maintaining the existing FC path information. Otherwise, if it is the source AS, the FC-BGP speaker generates the FC path attribute defined in Figure 1 and inserts it into the UPDATE message.

There are three AS numbers in one FC segment, as Figure 2 shows. The populating of the Current AS number (CASN) within the FC segment is like the AS number in BGPsec; it MUST match the AS number in the Subject field of the RPKI router certificate that will be used to verify the FC segment constructed by this FC-BGP speaker (see Section 3.1.1 of [RFC8209] and [RFC6487]). The Previous AS number (PASN) is typically set to the AS number from which the UPDATE message is received. However, if the FC-BGP speaker is located in the origin AS, the PASN SHOULD be filled with 0. The Nexthop AS number (NASN) is set to the AS number of the peer to whom the route is advertised. So if there are several neighbors, the FC-BGP speaker should generate separate FCs for different neighbors. But it would never generate a new FC segment for the iBGP neighbor.

The Subject Key Identifier field (SKI) within the new FC segment is populated with the identifier found in the Subject Key Identifier extension of the RPKI router certificate associated with the FC-BGP speaker. This identifier serves as a crucial piece of information for recipients of the route advertisement. It enables them to identify the appropriate certificate to employ when verifying the signatures in FC segments attached to the route advertisement. This practice adheres to the guidelines outlined in [RFC8209].

Typically, the Flags field is set to 0.

A route server (RS) is a third-party brokering system that interconnects three or more BGP-speaking routers using eBGP in IXPs [RFC7947]. Typically, RS performs like a transit AS except that it does not insert its AS number into the AS\_PATH attribute. The route server can also participate in the FC-BGP process. If the RS is an FC-BGP-enabled RS, it may choose to set the Flags-RS bit to 1 when it populates its FC segment. However, when the RS chooses to add its AS number to the AS\_PATH attribute, the Flags-RS bit SHOULD be set to 0. If the RS is a non-FC-BGP RS, it propagates the FC-BGP UPDATE message directly. Anyway, the AS number of RS would be used in the FC segment, no matter if it appears in the AS\_PATH attribute.

Typically, the route server does not insert ASN into AS\_PATH. Take Figure 3 as an example, where AS A advertises a BGP UPDATE to AS C and RS connects AS A and AS B. When RS supports the FC-BGP mechanism, AS A adds its FC segment: FC(NULL, A, RS), RS adds its FC segment: FC(A, RS, B, Flags-RS), and AS B adds its FC segment: FC(RS, B, C). If RS does not support the FC-BGP mechanism, FC(A, RS, B, Flags-RS) is missed in FCList, which SHOULD be considered as a partial deployment scenario.

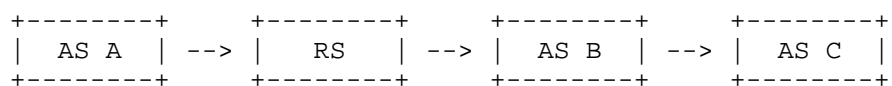


Figure 3: A network topology with a Router Server linking two ASs.

AS Path Prepending is a traffic engineering mechanism in BGP to deprioritize a route or alternate path, which will prepend the local AS number multiple times to the AS\_PATH attribute [ASPP]. To minimize unnecessary processing load during the validation of FC segments, an FC-BGP speaker SHOULD NOT generate multiple consecutive FC segments with the same AS number. Instead, the FC-BGP speaker SHOULD aim to produce a single FC segment once, even if the intention is to achieve the semantics of prepending the same AS number multiple times. Flags-ASPP MUST be set to 1 if the local AS uses AS Path Prepending.

The Algorithm ID field is set to 1. FC-BGP only supports one algorithm suite in this document as the BGPsec Algorithm defined in [RFC8208]. That is the signature algorithm used MUST be the Elliptic Curve Digital Signature Algorithm (ECDSA) with curve P-256, and the hash algorithm used MUST be SHA-256. If it transits from one algorithm suite to another, the FC-BGP speaker MUST place its router certificate in the RPKI repository first and then specify the Algorithm ID in the FC segment.

The Signature Length field is populated with the length (in octets) of the value in the Signature field.

The Signature field in the new FC segment is a variable-length field. It contains a digital signature encapsulated in DER format that binds the prefix, its length, and triplet <PASN, CASN, NASN> to the RPKI router certificate corresponding to the FC-BGP speaker. The digital signature is computed as follows: Signature=ECDSA(SHA256(PASN, CASN, NASN, SKI, Algorithm ID, Flags, Signature Length, Prefix, Prefix Length)).

The signatures within the FC segments of an FC-BGP UPDATE message ensure the protection of crucial information, including the AS number of the neighbor involved in the message exchange. This information is explicitly included in the generated FC segment. Consequently, if an FC-BGP speaker intends to transmit an FC-BGP UPDATE message to multiple BGP neighbors, it MUST generate a distinct FC-BGP UPDATE message for each unique neighbor AS to whom the UPDATE message is being sent.

Indeed, an FC-BGP UPDATE message is REQUIRED to advertise a route to just one prefix. This is because if an FC-BGP speaker receives an UPDATE message containing multiple prefixes, it would be unable to construct a valid FC-BGP UPDATE message, including valid path signatures, with a subset of the received prefixes. To advertise routes to multiple prefixes, the FC-BGP speaker MUST generate individual FC-BGP UPDATE messages for each prefix. This ensures the proper construction of valid path signatures for each advertised prefix.

All FC-BGP UPDATE messages MUST conform to BGP's maximum message size. If the resulting message exceeds the maximum message size, then the guidelines in Section 9.2 of [RFC4271] MUST be followed.

## 4.2. Propagation

Any BGP speaker should propagate the optional transitive FC path attribute encapsulated in the FC-BGP UPDATE message, even though they do not support the FC-BGP feature. However, it is important to note that an FC-BGP speaker SHOULD NOT make any attestation regarding the validation state of the FC-BGP UPDATE message it receives. They MUST verify the FC path attribute themselves.

To incorporate or create a new FC segment for an FC-BGP UPDATE message using a specific algorithm suite, the FC-BGP speaker MUST possess an appropriate private key capable of generating signatures for that particular algorithm suite. Moreover, this private key MUST correspond to the public key found in a valid RPKI End Entity (EE) certificate. The AS number resource extension within this certificate should include the FC-BGP speaker's AS number as specified in [RFC8209]. It is worth noting that these new segments are only prepended to an FC-BGP UPDATE message when the FC-BGP speaker generates the UPDATE message for transmission to an external neighbor. In other words, this occurs when the AS number of the neighbor differs from the AS number of the FC-BGP speaker.

The RPKI allows the legitimate holder of IP address prefix(es) to issue a digitally signed object known as a Route Origin Authorization (ROA). This ROA authorizes a specific AS to originate routes for a particular set of prefixes [RFC6482]. It is anticipated that most Relying Parties (RPs) will combine FC-BGP with origin validation [RFC6483] and [RFC6811]. Therefore, it is strongly RECOMMENDED that an FC-BGP speaker only advertises a route for a given prefix in an FC-BGP UPDATE message if there is a valid ROA that authorizes the FC-BGP speaker's AS to originate routes for that specific prefix.

If an FC-BGP router receives a non-FC-BGP UPDATE message from an external neighbor, meaning that the origin BGP speaker does not support FC-BGP, the router processes the UPDATE message as a regular BGP UPDATE message typically. In this case, it SHOULD NOT add its own FC segment to the UPDATE message. On the other hand, when the FC-BGP speaker advertises routes belonging to its local AS and receives them from an internal neighbor, it MUST add its FC segment to the UPDATE message if it decides to propagate those routes. Furthermore, if the FC-BGP router receives an FC-BGP UPDATE message from a neighbor for a specific prefix and chooses to propagate that neighbor's route for the prefix, it MUST propagate the route as an FC-BGP UPDATE message containing the FC path attribute.

When an FC-BGP speaker sends an FC-BGP UPDATE message to an iBGP (internal BGP) neighbor, the process is straightforward. When the FC-BGP speaker originates a new route advertisement and sends it to

an iBGP neighbor, it MUST NOT include the FC path attribute of the UPDATE message. In other words, the FC-BGP speaker omits the FC path attribute. Similarly, when an FC-BGP speaker decides to forward an FC-BGP UPDATE message to an iBGP neighbor, it MUST refrain from adding a new FC segment to the FC-BGP UPDATE message.

When an FC-BGP speaker receives an FC-BGP UPDATE message containing an FC path attribute (with one or more FC segments) from an (internal or external) neighbor, it may choose to propagate the route advertisement by sending it to its other (internal or external) neighbors. When sending the route advertisement to an internal FC-BGP-speaking neighbor, the FC path attribute SHALL NOT be modified. When sending the route advertisement to an external FC-BGP-speaking neighbor, the following procedures are used to form or update the FC path attribute.

#### 4.3. Processing Instructions for AS Confederation Members

Members of AS Confederation [RFC5065] MUST additionally follow the instructions in this section for processing FC-BGP UPDATE messages.

When an FC-BGP speaker in an AS confederation receives an FC-BGP UPDATE message from a neighbor that is external to the confederation and chooses to propagate the UPDATE message within the confederation, it first adds an FC segment and the signature signed to its own Member-AS (i.e., the 'Current AS Number' is the FC-BGP speaker's Member-AS Number). In this internally modified UPDATE message, the newly added FC segment contains the public AS number (i.e., Confederation Identifier), and the segment's `Confed_Segment` flag is set to 1. The newly added signature is generated using a private key corresponding to the public AS number of the confederation. The FC-BGP speaker propagates the modified UPDATE message to its peers within the confederation. (Note: In this document, intra-Member-AS peering is regarded as iBGP, and inter-Member-AS peering is regarded as eBGP. The latter is also known as confederation-eBGP.)

Within a confederation, the verification of FC-BGP signatures added by other members of the confederation is optional. Note that if a confederation chooses not to verify digital signatures within the confederation, then FC-BGP is not able to provide any assurances about the integrity of the Member-AS Numbers placed in FC segments where the `Confed_Segment` flag is set to 1.

When a confederation member receives an FC-BGP UPDATE message from a peer within the confederation and propagates it to a peer outside the confederation, it needs to remove all of the FC segments added by confederation members when it removes all path segments of the AS\_PATH with the type of AS\_CONFED\_SEQUENCE or AS\_CONFED\_SET. To do this, the confederation members propagated the route outside the confederation as follows:

1. Starting with the most recently added FC segment, remove FC segments whose Flags-CS bit is 1. Stop this process once an FC segment that has its Confed\_Segment flags set to 0 is reached.
2. Add an FC segment containing, in the CASN field, the AS Confederation Identifier (the public AS number of the confederation). Note that all fields other than the CASN field are populated.

Finally, as discussed above, an AS confederation MAY optionally decide that its members will not verify digital signatures added by members. In such a confederation, when an FC-BGP speaker runs the algorithm in Section 5.2.1, the FC-BGP speaker, during the process of signature verifications, first checks whether the Confed\_Segment flag in an FC segment is set to 1. If the flag is set to 1, the FC-BGP speaker skips verification for the corresponding signature and proceeds immediately to the next FC segment. It is an error when an FC-BGP speaker receives, from a neighbor who is not in the same AS confederation, an FC-BGP UPDATE message containing a Confed\_Segment flag set to 1.

#### 4.4. Processing Instructions for BGP Route Leak Prevention

The BGP routing system is susceptible to numerous vulnerabilities. The systemic vulnerability of the BGP routing system is known as "route leaks" [RFC7908]. There are 6 types of route leaks defined in [RFC7908].

[RFC9234] can detect and prevent BGP route leaks by adding a new BGP OPEN Role capability and OTC transitive path attribute. However, it may be forged.

When using BGP route leak prevention with FC-BGP, it SHOULD tell the neighbor its BGP Role as Section 4 of [RFC9234]. If the peer's role is Customer or RS-Client, the Flags-P2C MUST be set to 1 on route advertisement. If the peer's role is peer, the Flags-P2P MUST be set to 1 on the route advertisement. Then, the route SHOULD subsequently go only to the Customers. It is worth noting that if the recently added FC Segment already set the Flags-P2C or Flags-P2P flags to 1, it MUST NOT be propagated to Providers, Peers, or RSes.

## 5. Processing a Received FC-BGP UPDATE Message

### 5.1. Overview

When receiving an FC-BGP UPDATE message from an external BGP neighbor carrying the FC path attribute, an FC-BGP speaker SHOULD first validate the message to determine the authenticity of the path information. Same as BGPsec, an FC-BGP speaker will wish to perform origin validation (see [RFC6483] and [RFC6811]) on an incoming FC-BGP UPDATE message, but such validation is independent of the validation described in this section.

After the validation, the FC-BGP speaker may want to send the FC-BGP UPDATE message to neighbors according to local route policies. Then it SHOULD update the FC path attributes and continue advertising the BGP route.

For the origin AS that launches the advertisement, the FC-BGP speaker only needs to generate the FC-BGP UPDATE message without the validation.

The FC-BGP speaker stores the router certificates in the RPKI repository, and any changes in the RPKI state can impact the validity of the UPDATE messages. That means the validity of FC-BGP UPDATE messages relies on the current state of the RPKI repository. When an FC-BGP speaker becomes aware of a change in the RPKI state, such as through an RPKI validating cache using the RTR protocol (as specified in [RFC8210]), it is REQUIRED to rerun validation on all affected UPDATE messages stored in its Adj-RIB-In [RFC4271]. For instance, if a specific RPKI router certificate becomes invalid due to expiration or revocation, all FC-BGP UPDATE messages containing an FC segment with an SKI matching the SKI in the affected certificate must be reassessed to determine their current validity. If the reassessment reveals a change in the validity state of an UPDATE message, the FC-BGP speaker, depending on its local policy, SHOULD rerun the best path selection process. This allows for the appropriate handling of the updated information and ensures that the most valid and suitable paths are chosen for routing purposes.

### 5.2. Validation

When verifying the authenticity of an FC-BGP UPDATE message, information from the RPKI router certificates is utilized. The RPKI router certificates provide the data, including the triplet <AS Number, Public Key, Subject Key Identifier>, to verify the AS\_PATH and FC path attributes. As a prerequisite, the recipient MUST have access to these RPKI router certificates.

Note that if an AS uploads its router certificates to RPKI, it would be deemed to support FC-BGP. The validation process MUST check to ensure no malicious on-path AS removes FCs from FC-BGP UPDATE.

Note that the FC-BGP speaker could perform the validation of RPKI router certificates on its own and extract the required data, or it could receive the same data from a trusted cache that performs RPKI validation on behalf of (some set of) FC-BGP speakers. (For example, the trusted cache could deliver the necessary validity information to the FC-BGP speaker by using the Router Key PDU (Protocol Data Unit) for the RPKI-Router protocol [RFC8210].)

The recipient validates an FC-BGP UPDATE message containing the FC path attribute and obtains one of two states: 'Valid' and 'Not Valid'. We will describe the validation procedure in Section 5.2.1 in this document. The validation result will be used at BGP route selection, thus it will be discussed at Section 6.4.

As the FC-BGP UPDATE message is generated at the eBGP router, the FC-BGP validation needs only to be performed at the eBGP router. The iBGP route plays a crucial role in the FC-BGP UPDATE message propagation and distribution. The function of iBGP is to convey the validation status of an FC-BGP UPDATE message from an ingress edge router to an egress edge router within an AS. The specific mechanisms used to convey the validation status can vary depending on the implementation and local policies of the AS. By propagating this information through iBGP, the eBGP router and other routers within the AS can be aware of the validation status of the FC-BGP UPDATE messages and make routing decisions accordingly. As stated in Section 4, when an FC-BGP speaker decides to forward a syntactically correct FC-BGP UPDATE message, it is RECOMMENDED to do so while preserving the FC path attribute. This recommendation applies regardless of the validation state of the UPDATE message.

Ultimately, the decision to forward the FC-BGP UPDATE message with the FC path intact and the choice to perform independent validation at the egress router are both determined by local policies implemented within the AS. Note that the decision to perform validation on the received FC-BGP UPDATE message is left to the discretion of the egress router, which is the router receiving the message within its own AS. The egress router has the freedom to choose whether or not it wants to independently validate the FC path attribute based on its local policy, even if the FC path attribute has already been validated by the ingress router. This additional validation performed at the egress router helps ensure the integrity and security of the received FC-BGP UPDATE message.

### 5.2.1. Validation Algorithm

This section specifies the concrete validation algorithm of FC-BGP UPDATE messages. A compliant implementation MUST have an FC-BGP UPDATE validation algorithm that behaves the same as the specified algorithm. This ensures consistency and security in validating FC-BGP UPDATE messages across different implementations. It allows for interoperability and standardized communication between FC-BGP-enabled networks.

First, the integrity of the FC-BGP UPDATE message MUST be checked. Both syntactical and protocol violation errors are checked. The FC path attribute MUST be present when an FC-BGP UPDATE message is received from an external FC-BGP neighbor and also when such an UPDATE message is propagated to an internal FC-BGP neighbor. The error checks specified in Section 6.3 of [RFC4271] are performed, except that for FC-BGP UPDATE messages, the checks on the FC path attribute do not apply, and instead, the following checks on the FC path attribute are performed:

1. Check to ensure that the entire FC path attribute is syntactically correct (conforms to the specification in this document).
2. For each AS on AS\_PATH that has a router certificate in RPKI, check whether the FC path attribute contains a corresponding FC segment whose CASN field has the same value as the AS number.
3. Check that the triplet <PASN, CASN, NASN> fields in each FC segment follow the order in AS\_PATH.
4. Check that each FC segment contains one signature with the supported Algorithm ID.
5. If the UPDATE message was received from an FC-BGP neighbor that is not a member of the FC-BGP speaker's AS confederation, check to ensure that none of the FC Segments contain a Flags field with the Confed\_Segment flag set to 1.
6. If the UPDATE message was received from an FC-BGP neighbor that is not a member of the FC-BGP speaker's AS confederation, check to ensure that the FC Segment corresponding to that peer does not contain a Flags field with the Flags-CS flag set to 1.
7. If the UPDATE message was received from an FC-BGP neighbor that is a member of the FC-BGP speaker's AS confederation, check to ensure that the FC Segment corresponding to that peer contains a Flags field with the Flags-CS flag set to 1.

8. If the UPDATE message was received from a neighbor that is not expected to set Flags-RS bit to 0 (see Section 4), then check to ensure that the Flags-RS bit in the most recently added FC Segment is not equal to 0.
9. If the UPDATE message was received from a neighbor that is not expected to set the Flags-ASPP bit to 0, i.e., the AS doesn't use the ASPP mechanism, then check to ensure that the Flags-RS bit in the most recently added FC Segment is not equal to 0.
10. If the UPDATE message was received from a neighbor that is expected to set the Flags-RS bit to 0 (see Section 4), then check to ensure that the Flags-RS bit in the most recently added FC Segment is equal to 0.
11. If the UPDATE message was received from a neighbor that is not expected to set the Flags-P2C bit or Flags-P2P bit to 1 (see Section 4), then check to ensure that the Flags-P2C bit or Flags-P2P bit in the most recently added FC Segment is not equal to 1.
12. If the UPDATE message was received from a neighbor that is expected to set the Flags-P2C bit or Flags-P2P bit to 1 (see Section 4), then check to ensure that the Flags-P2C bit or Flags-P2P bit in the most recently added FC Segment is equal to 1.

If any of the checks for the FC path attribute fail, indicating a syntactical or protocol error, it is considered an error. In such cases, FC speakers are REQUIRED to handle these errors using the "treat-as-withdraw" approach as defined in [RFC7606]. This approach means that the FC-BGP speaker SHOULD treat the FC path attribute as if it were a withdraw message, effectively removing the route from consideration. It's worth noting that when a transparent route server is involved, and its AS number appears in the FC (with the Flags-RS bit set to 1), the route server has the option to check if its local AS is listed in the FC. This additional check can be included as part of the loop-detection mechanism mentioned earlier in the specification.

When an AS appears on the AS\_PATH of an UPDATE message and has uploaded router certificates in RPKI, it MUST add its FC segment to the FC path attribute. Otherwise, the downstream ASs SHOULD consider that the FC of this AS has been removed by other ASs and the UPDATE message is falsified.

When one FC Segment has set the Flags-P2C flag to 1, the subsequent FC Segments added by the following ASs MUST all set the Flags-P2C flag to 1 in their corresponding FC Segments. The Flags-P2C flag is set to 1 only when the role of its neighbor, to whom the propagator AS sends routes, is Customer or RS-Client. The Flags-P2P flag is set to 1 only when the role of its neighbor, to whom the propagator AS sends routes, is Peer.

The following ingress procedure applies to the processing of the Flags-P2C flag and the Flags-P2P flag on route receipt:

1. If a route, with the Flags-P2C flag in the recently added FC Segment, is received from a Customer, a Peer, or an RS-Client, then it is a route leak and MUST be considered ineligible.
2. If a route is received from a Peer (i.e., remote AS with a Peer Role) and with the Flags-P2P flag in both two recently consecutively added FC Segments, then it is a route leak and MUST be considered ineligible.

Next, the FC-BGP speaker iterates through the FC segments. Once the FC-BGP speaker has examined the signature field in the FC attribute, it proceeds to validate the signature using the supported algorithm suites. However, if the FC-BGP speaker encounters a signature corresponding to an algorithm suite indexed by an Algorithm ID that it does not support, that particular signature is not considered in the validation process. If there are no signatures corresponding to any algorithm suites supported by the FC-BGP speaker, a specific action is taken to ensure the continuity of the route selection process. To consider the UPDATE message in the route selection process, the FC-BGP speaker has to treat the message as if it were received as an unsigned BGP UPDATE message. By treating the UPDATE message as unsigned, the FC-BGP speaker acknowledges that it cannot verify the integrity and authenticity of the message through the provided signatures. However, it still allows the message to be considered for route selection, ensuring that important routing information is not disregarded solely due to the lack of supported signature algorithms.

For each remaining signature corresponding to an algorithm suite supported by the FC-BGP speaker, the FC-BGP speaker processes FC-BGP UPDATE message validation with the following steps. As different FC segments are independent, it is recommended to verify FC segments parallelly.

- \* Step 1: Locate the public key needed to verify the signature in the current FC segment. To do this, consult the valid RPKI router certificate data and look up all valid <AS Number, Public Key,

Subject Key Identifier> triples in which the AS matches the Current AS Number (CASN) in the corresponding FC segment. Of these triples that match the AS number, check whether there is an SKI that matches the value in the Subject Key Identifier field of the FC segment. If this check finds no such matching SKI value, then mark the entire FC segment as 'Not Valid' and stop.

- \* Step 2: Compute the digest function (for the given algorithm suite) on the appropriate data. To verify the digital signature in the FC segment, construct the sequence of octets to be hashed. Note that this sequence is the same sequence that was used by AS that created the FC Segment (see Section 4). The elements in this sequence MUST be ordered exactly as shown in the generation process. Note that if an FC-BGP speaker uses multiple AS Numbers (e.g., the FC-BGP speaker is a member of a confederation), the AS number used here MUST be the AS number announced in the OPEN message for the BGP session over which the FC-BGP UPDATE message was received. All three AS numbers in one FC segment follow this rule.
- \* Step 3: Use the signature validation algorithm (for the given algorithm suite) to verify the signature in the current segment. That is, invoke the signature validation algorithm on the following three inputs: the value of the signature field in the current FC segment, the digest value computed in Step 2 above, and the public key obtained from the valid RPKI data in Step 1 above. If the signature validation algorithm determines that the signature is invalid, then mark the entire FC segment as 'Not Valid' and stop. If the signature validation algorithm determines that the signature is valid, then the FC segment is marked as 'Valid' and continues to process the following FC segments.

If all FC segments are marked as 'Valid', then the validation algorithm terminates, and the FC-BGP UPDATE message is deemed 'Valid'. Otherwise, the FC-BGP UPDATE message is deemed 'Not Valid'.

## 6. Implementations, Operations, and Management Considerations

### 6.1. Algorithms and Extensibility

The content of Algorithm Suite Considerations defined in Section 6.1 of [RFC8205] and the content of Considerations for the SKI Size defined in Section 6.2 of [RFC8205] are indeed applicable in this context of FC-BGP.

But the algorithm suite transition in FC-BGP is straightforward: As each FC segment has an Algorithm ID field, just populate this field with a feasible and consensus value that all FC-BGP speaker supports when transitioning.

## 6.2. Speedup and Early Termination of Signature Verification

It is advantageous for an implementation to establish a parallel verification process for FC-BGP if the router's processor supports such operations. As each FC segment contains the integral data that needs to be verified, parallel verification can significantly enhance the efficiency and speed of the validation process. By utilizing parallel processing capabilities, an implementation can simultaneously verify multiple FC segments, thereby reducing the overall verification time. This is particularly beneficial in scenarios where the FC path attribute contains a substantial number of segments or in high-traffic networks with a large volume of FC-BGP UPDATE messages. Implementations that leverage parallel verification take advantage of the processing power available in modern router processors. This allows for more efficient and faster verification, ensuring that the FC-BGP UPDATE messages are promptly validated and routed accordingly.

However, it's important to note that the feasibility of parallel verification depends on the specific capabilities and constraints of the router's processor. Implementations SHOULD consider factors such as available resources, concurrency limitations, and the impact on overall system performance when implementing parallel verification processes. Overall, setting up a parallel verification process for FC-BGP, if feasible, can contribute to improved performance and responsiveness in validating FC segments, further enhancing the reliability and efficiency of the FC-BGP protocol.

During the validation of an FC-BGP UPDATE message, route processor performance speedup can be achieved by incorporating the following observations. These observations provide valuable insights into optimizing the validation process and reducing the workload on the route processor. One of the key observations is that an FC-BGP UPDATE message is considered 'Valid' only if all FC segments are marked as 'Valid' in the validation steps. This means that if an FC segment is marked as 'Not Valid', there is no need to continue verifying the remaining unverified FC segments. This optimization can significantly reduce the processing time and workload on the route processor. Furthermore, when the FC-BGP UPDATE message is selected as the best path, the FC-BGP speaker appends its own FC segment, including the appropriate signature generated with the corresponding algorithm, to the FC path attribute. This ensures that the updated path is propagated correctly.

Additionally, an FC-BGP UPDATE message is considered as 'Not Valid' if at least one signature in each of the FC segments is invalid. Thus, the verification process for an FC segment can terminate early as soon as the first invalid signature is encountered. There is no need to continue validating the remaining signatures in that FC segment.

By incorporating these observations, an FC-BGP implementation can achieve significant performance improvements and reduce the computational burden on the route processor. It allows for more efficient validation of FC-BGP UPDATE messages, ensuring the integrity and security of the routing information while maximizing system resources.

### 6.3. Deferring Validation

When an FC-BGP speaker receives an exceptionally large number of UPDATE messages simultaneously, though it can use parallel verification to speed up the validation, it can be beneficial to defer the validation of incoming FC-BGP UPDATE messages. The decision to defer the validation process may depend on the local policy of the FC-BGP speaker, taking into account factors such as available resources and system load.

By deferring the validation of these messages, the FC-BGP speaker can prioritize its processing power and resources to handle other critical tasks or ongoing operations. Deferring the validation allows the FC-BGP speaker to temporarily postpone the resource-intensive validation process until it can allocate sufficient resources to handle the influx of incoming messages effectively.

The implementation SHOULD provide visibility to the operator regarding the deferment of validation and the status of the deferred messages. This visibility enables the operator to have awareness of the deferred messages and understand the current state of the system. This information is crucial for monitoring and managing the FC-BGP speaker's behavior, ensuring that the operator can make informed decisions based on the system's status.

### 6.4. BGP Route Selection

While FC-BGP does modify the BGP route selection result, it is not the primary intention of FC-BGP to modify the BGP route selection process itself. Instead, FC-BGP focuses on providing an additional layer of validation and verification for BGP UPDATE messages.

However, the handling of FC-BGP validation states, as well as the integration of FC-BGP with the BGP route selection, is indeed a matter of local policy. FC-BGP implementations SHOULD provide mechanisms that allow operators to define and configure their own local policies on a per-session basis. This flexibility enables operators to customize the behavior of FC-BGP based on their specific requirements and preferences.

By allowing operators to set local policies, FC-BGP implementations empower them to control how the validation status of FC-BGP UPDATE messages influences the BGP route selection process. Operators may choose to treat FC-BGP validation status differently for UPDATE messages received over different BGP sessions, based on their network's needs and security considerations.

To ensure consistency and interoperability, it is RECOMMENDED that FC-BGP implementations treat the priority of FC-BGP UPDATE messages at the same level as Route Origin Validation (ROV). This means that the validation status of FC-BGP UPDATE messages should be considered alongside other route selection criteria, such as path attributes, AS path length, and local preference.

#### 6.5. Non-deterministic Signature Algorithms

The non-deterministic nature of many signature algorithms can introduce variations in the signatures produced, even when signing the same data with the same key. This means that if an FC-BGP router receives two FC-BGP UPDATE messages from the same peer, for the same prefix, with the same FC path attribute except the signature fields, the signature fields MAY differ when using a non-deterministic signature algorithm. Note that if the sender caches and reuses the previous signature, the two sets of signature fields will not differ. This applies specifically to deterministic signature algorithms, where the signature fields between the two UPDATE messages MUST be identical.

Considering these observations, an FC-BGP implementation MAY incorporate optimizations in the UPDATE validation processing. These optimizations can take advantage of the non-deterministic nature of signature algorithms to reduce computational overhead. For example, if an FC-BGP router has already validated an FC segment and its corresponding signature in a previous UPDATE message from the same peer, it may choose to cache and reuse the previous validation result. This can help avoid redundant computations for subsequent UPDATE messages with the same FC path attribute and SKIs, as long as the sender does not generate new signatures.

By incorporating such optimizations, an implementation can reduce the computational load and processing time needed for validating FC-BGP UPDATE messages. However, it is important to ensure that the implementation adheres to the requirements and specifications of the FC-BGP protocol while considering the performance benefits of these optimizations.

#### 6.6. Private AS Numbers

The process of Private AS Numbers used in BGPsec speaker defined in Section 7.5 of [RFC8205] also applies here.

#### 6.7. Robustness Considerations for Accessing RPKI Data

As there is a mature RPKI to Router protocol [RFC8210], the implementation is REQUIRED to use this protocol to access the RPKI data. The content defined in Section 7.6 of [RFC8205] also applies here.

#### 6.8. Graceful Restart

During Graceful Restart (GR), restarting and receiving FC-BGP speakers MUST follow the procedures specified in [RFC4724] for restarting and receiving BGP speakers, respectively. In particular, the behavior of retaining the forwarding state for the routes in the Loc-RIB [RFC4271] and marking them as stale, as well as not differentiating between stale routing information and other information during forwarding, will be the same as the behavior specified in [RFC4724].

#### 6.9. Robustness of Secret Random Number in ECDSA

As both FC-BGP and BGPsec use ECDSA, the content of Robustness of Secret Random Number in ECDSA defined in Section 7.8 of [RFC8205] applies here.

#### 6.10. Incremental/Partial Deployment Considerations

The core difference between FC-BGP and BGPsec is that BGPsec is a path-level authentication approach, whereas FC-BGP is a pathlet-driven authentication approach.

In design, FC-BGP does not modify the AS\_PATH attribute. It defines a new transitive path attribute to transport the FC segments so that the legacy ASs can forward this attribute to their peers. Thus, FC-BGP is natively compatible with BGP and supports partial deployment. It differs from BGPsec, which replaces the AS\_PATH attribute with a new Secure\_Path information of the BGPsec\_Path attribute.

As for incremental/partial deployment considerations, in Section 5.1.1 of [FC-ARXIV], we have proved that the adversary cannot forge a valid AS path when FC-BGP is universally deployed. Section 5.1.2 of [FC-ARXIV] analyzes the benefits of FC-BGP in the case of partial deployment. The results show that FC-BGP provides more benefits than BGPsec in partial deployment. As a result, attackers are forced to pretend to be at least two hops away from the destination AS, which reduces the probability of successful path hijacks.

#### 6.11. Co-existence with BGPsec

It is NOT RECOMMENDED that both BGPsec and FC-BGP be enabled together. However, at the very least, the implementation SHOULD adequately process the coexistence update message.

When an FC-BGP speaker also enables the BGPsec feature, it MUST also properly process the BGPsec UPDATE message as follows:

- \* General Principle. The BGP speaker SHOULD prioritize BGPsec over FC-BGP. When both features are enabled, the BGP speaker processes the BGPsec UPDATE message first, then processes the FC-BGP UPDATE message.
- \* FC-BGP UPDATE Message Generation. The BGP speaker SHOULD prioritize the BGPsec\_Path attribute over the AS\_PATH attribute. This means that when broadcasting a BGP UPDATE message, the BGP speaker SHOULD first check if the peer supports BGPsec. If so, it SHOULD generate a BGPsec-enabled UPDATE message. In this message, BGPsec\_Path replaces the AS\_PATH attribute, and an MP\_REACH\_NLRI attribute [RFC4760] is used for encoding NLRI information. The generation of the FC Path attribute SHOULD use these attributes to generate FC Segments. The impact on FC generation is minimal, as it only needs to obtain PASN, CASN, NASN, Prefix Address, and Prefix Length from BGPsec\_Path and MP\_REACH\_NLRI attributes separately.
- \* FC-BGP UPDATE Message Validation. The BGP speaker should also prioritize the BGPsec\_Path over AS\_PATH. After processing the BGPsec path attribute, the BGP speaker should decode the BGPsec\_Path and MP\_REACH\_NLRI attributes. So, in the validation process, the BGP speaker SHOULD essentially reverse the steps it took during generation. It would first decode the BGPsec\_Path and MP\_REACH\_NLRI attributes to obtain the necessary information (PASN, CASN, NASN, Prefix Address, and Prefix Length). Then, it would use this information to validate the corresponding FC Segment.

In summary, the coexistence of BGPsec and FC-BGP is not overly burdensome.

## 7. Security Considerations

### 7.1. Security Guarantees

TBD.

When FC-BGP is used in conjunction with origin validation, the following security guarantees can be achieved:

- \* The source AS in a route announcement is authorized.
- \* FC-BGP speakers on the AS-Path are authorized to propagate the route announcements.
- \* The forwarding path of packets is consistent with the routing path announced by the FC-BGP speakers.

FC-BGP is designed to enhance the security of control plane routing in the Internet at the network layer. Specifically, FC-BGP allows an AS to independently prove its BGP routing decisions with publicly verifiable cryptography commitments, based on which any on-path AS can verify the authenticity of a BGP path. More crucially, the above security guarantees offered by FC-BGP are not binary, i.e., secure or non-secure. Instead, the security benefits are strictly monotonically increasing as the deployment rate of FC-BGP (i.e., the percentage of ASs that are upgraded to support FC-BGP) increases.

### 7.2. Mitigation of Denial-of-Service Attacks

The FC-BGP UPDATE process, due to its involvement in numerous cryptographic operations, becomes vulnerable to Denial-of-Service (DoS) attacks targeting FC-BGP speakers. This section addresses the mitigation strategies tailored for the specific DoS threats the FC-BGP protocol poses. To prevent the Denial-of-Service (DoS) attacks faced by the FC-BGP control plane mechanism, there is no need to put in more effort than BGPsec.

To reduce the impact of DoS attacks, FC-BGP speakers SHOULD employ an UPDATE validation algorithm that prioritizes inexpensive checks (such as syntax checks) before proceeding to more resource-intensive operations (like signature verification). The validation algorithm described in Section 5.2.1 is designed to sequence checks in order of likely expense, starting with less costly operations. However, the actual cost of executing these validation steps can vary across different implementations, and the algorithm in Section 5.2.1 may not offer the optimal level of DoS protection for all cases.

Moreover, the transmission of UPDATE messages with the FC path attribute, which entails a multitude of signatures, is a potential vector for denial-of-service attacks. To counter this, implementations of the validation algorithm must cease signature verification immediately upon encountering an invalid signature. This prevents prolonged sequences of invalid signatures from being exploited for DoS purposes. Additionally, implementations can further mitigate such attacks by limiting validation efforts to only those UPDATE messages that, if found to be valid, would be chosen as the best path. In other words, if an UPDATE message includes a route that would be disqualified by the best path selection process for some reason (such as an excessively long AS path), it is OPTIONALLY to determine its FC-BGP validity status.

### 7.3. Route Server

When the Route Server populates its FC Segment into the FC path attribute, it is secure as the path is fully deployed.

When the Route Server fails to insert the FC Segment, no matter whether its ASN is listed in the AS path, it is considered a partial deployment, which poses a risk of path forgery.

### 7.4. Additional Security Considerations

#### 7.4.1. Three AS Numbers

An FC segment contains only partial path information, and FCs in the FCList are independent. To prevent BGP Path Splicing attacks, we propose to use the triplet <Previous AS Number, Current AS Number, Nexthop AS Number> to locate the pathlet information.

But if there is no previous hop, i.e., this is the origin AS that tries to add its FC segment to the BGP UPDATE message, the Previous AS Number SHOULD be populated with 0. But, carefully, AS 0 SHOULD only be used in this case.

In the context of BGP [RFC4271], to detect an AS routing loop, it scans the full AS path (as specified in the AS\_PATH attribute) and checks that the autonomous system number of the local system does not appear in the AS path. As outlined in [RFC7607], Autonomous System 0 was listed in the IANA Autonomous System Number Registry as "Reserved - May be used to identify non-routed networks". So, there should be no AS 0 in the AS\_PATH attribute of the BGP UPDATE message. Therefore, AS 0 could be used to populate the PASN field when there are no previous AS hops in the AS path.

#### 7.4.2. MISC

For a discussion of the BGPsec threat model and related security considerations, please see [RFC7132]. The security considerations of [RFC4272] also apply to FC-BGP.

### 8. IANA Considerations

TBD. Wait for IANA to assign FC-BGP-UPDATE-PATH-ATTRIBUTE-TYPE.

TBD. Regist Flags. The leftmost bit is the Confed\_Segment flag, and the second highest/leftmost bit is the Route\_Server flag in this document.

TBD. A new OID should be assigned for keys used in FC-BGP.

AS number 0 is used here to populate the PASN in an FC segment where there is no previous hop for an AS, i.e., the origin AS when adding the FC segment to the FC-BGP UPDATE message.

### 9. References

#### 9.1. Normative References

- [RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <<https://www.rfc-editor.org/rfc/rfc4271>>.
- [RFC4724] Sangli, S., Chen, E., Fernando, R., Scudder, J., and Y. Rekhter, "Graceful Restart Mechanism for BGP", RFC 4724, DOI 10.17487/RFC4724, January 2007, <<https://www.rfc-editor.org/rfc/rfc4724>>.
- [RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, DOI 10.17487/RFC4760, January 2007, <<https://www.rfc-editor.org/rfc/rfc4760>>.

- [RFC5656] Stebila, D. and J. Green, "Elliptic Curve Algorithm Integration in the Secure Shell Transport Layer", RFC 5656, DOI 10.17487/RFC5656, December 2009, <<https://www.rfc-editor.org/rfc/rfc5656>>.
- [RFC6480] Lepinski, M. and S. Kent, "An Infrastructure to Support Secure Internet Routing", RFC 6480, DOI 10.17487/RFC6480, February 2012, <<https://www.rfc-editor.org/rfc/rfc6480>>.
- [RFC6482] Lepinski, M., Kent, S., and D. Kong, "A Profile for Route Origin Authorizations (ROAs)", RFC 6482, DOI 10.17487/RFC6482, February 2012, <<https://www.rfc-editor.org/rfc/rfc6482>>.
- [RFC6483] Huston, G. and G. Michaelson, "Validation of Route Origination Using the Resource Certificate Public Key Infrastructure (PKI) and Route Origin Authorizations (ROAs)", RFC 6483, DOI 10.17487/RFC6483, February 2012, <<https://www.rfc-editor.org/rfc/rfc6483>>.
- [RFC6487] Huston, G., Michaelson, G., and R. Loomans, "A Profile for X.509 PKIX Resource Certificates", RFC 6487, DOI 10.17487/RFC6487, February 2012, <<https://www.rfc-editor.org/rfc/rfc6487>>.
- [RFC6793] Vohra, Q. and E. Chen, "BGP Support for Four-Octet Autonomous System (AS) Number Space", RFC 6793, DOI 10.17487/RFC6793, December 2012, <<https://www.rfc-editor.org/rfc/rfc6793>>.
- [RFC7606] Chen, E., Ed., Scudder, J., Ed., Mohapatra, P., and K. Patel, "Revised Error Handling for BGP UPDATE Messages", RFC 7606, DOI 10.17487/RFC7606, August 2015, <<https://www.rfc-editor.org/rfc/rfc7606>>.
- [RFC7947] Jasinska, E., Hilliard, N., Raszuk, R., and N. Bakker, "Internet Exchange BGP Route Server", RFC 7947, DOI 10.17487/RFC7947, September 2016, <<https://www.rfc-editor.org/rfc/rfc7947>>.
- [RFC8205] Lepinski, M., Ed. and K. Sriram, Ed., "BGPsec Protocol Specification", RFC 8205, DOI 10.17487/RFC8205, September 2017, <<https://www.rfc-editor.org/rfc/rfc8205>>.
- [RFC8208] Turner, S. and O. Borchert, "BGPsec Algorithms, Key Formats, and Signature Formats", RFC 8208, DOI 10.17487/RFC8208, September 2017, <<https://www.rfc-editor.org/rfc/rfc8208>>.

- [RFC8209] Reynolds, M., Turner, S., and S. Kent, "A Profile for BGPsec Router Certificates, Certificate Revocation Lists, and Certification Requests", RFC 8209, DOI 10.17487/RFC8209, September 2017, <<https://www.rfc-editor.org/rfc/rfc8209>>.
- [RFC8210] Bush, R. and R. Austein, "The Resource Public Key Infrastructure (RPKI) to Router Protocol, Version 1", RFC 8210, DOI 10.17487/RFC8210, September 2017, <<https://www.rfc-editor.org/rfc/rfc8210>>.
- [RFC8635] Bush, R., Turner, S., and K. Patel, "Router Keying for BGPsec", RFC 8635, DOI 10.17487/RFC8635, August 2019, <<https://www.rfc-editor.org/rfc/rfc8635>>.
- [RFC9234] Azimov, A., Bogomazov, E., Bush, R., Patel, K., and K. Sriram, "Route Leak Prevention and Detection Using Roles in UPDATE and OPEN Messages", RFC 9234, DOI 10.17487/RFC9234, May 2022, <<https://www.rfc-editor.org/rfc/rfc9234>>.
- [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

- [RFC4272] Murphy, S., "BGP Security Vulnerabilities Analysis", RFC 4272, DOI 10.17487/RFC4272, January 2006, <<https://www.rfc-editor.org/rfc/rfc4272>>.
- [RFC5065] Traina, P., McPherson, D., and J. Scudder, "Autonomous System Confederations for BGP", RFC 5065, DOI 10.17487/RFC5065, August 2007, <<https://www.rfc-editor.org/rfc/rfc5065>>.
- [RFC5492] Scudder, J. and R. Chandra, "Capabilities Advertisement with BGP-4", RFC 5492, DOI 10.17487/RFC5492, February 2009, <<https://www.rfc-editor.org/rfc/rfc5492>>.

- [RFC6472] Kumari, W. and K. Sriram, "Recommendation for Not Using AS\_SET and AS\_CONFED\_SET in BGP", BCP 172, RFC 6472, DOI 10.17487/RFC6472, December 2011, <<https://www.rfc-editor.org/rfc/rfc6472>>.
- [RFC6811] Mohapatra, P., Scudder, J., Ward, D., Bush, R., and R. Austein, "BGP Prefix Origin Validation", RFC 6811, DOI 10.17487/RFC6811, January 2013, <<https://www.rfc-editor.org/rfc/rfc6811>>.
- [RFC7132] Kent, S. and A. Chi, "Threat Model for BGP Path Security", RFC 7132, DOI 10.17487/RFC7132, February 2014, <<https://www.rfc-editor.org/rfc/rfc7132>>.
- [RFC7607] Kumari, W., Bush, R., Schiller, H., and K. Patel, "Codification of AS 0 Processing", RFC 7607, DOI 10.17487/RFC7607, August 2015, <<https://www.rfc-editor.org/rfc/rfc7607>>.
- [RFC7908] Sriram, K., Montgomery, D., McPherson, D., Osterweil, E., and B. Dickson, "Problem Definition and Classification of BGP Route Leaks", RFC 7908, DOI 10.17487/RFC7908, June 2016, <<https://www.rfc-editor.org/rfc/rfc7908>>.
- [RFC8416] Ma, D., Mandelberg, D., and T. Bruijnzeels, "Simplified Local Internet Number Resource Management with the RPKI (SLURM)", RFC 8416, DOI 10.17487/RFC8416, August 2018, <<https://www.rfc-editor.org/rfc/rfc8416>>.
- [ASPP] McBride, M., Madory, D., Tantsura, J., Raszuk, R., Li, H., Heitz, J., and G. S. Mishra, "AS Path Prepending", Work in Progress, Internet-Draft, draft-ietf-grow-as-path-prepend-17, 17 October 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-grow-as-path-prepend-17>>.
- [ASPA-Profile] Snijders, J., Azimov, A., Uskov, E., Bush, R., Housley, R., and B. Maddison, "A Profile for Autonomous System Provider Authorization", Work in Progress, Internet-Draft, draft-ietf-sidrops-asma-profile-23, 24 March 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-sidrops-asma-profile-23>>.
- [ASPA-Verification] Azimov, A., Bogomazov, E., Bush, R., Patel, K., Snijders, J., and K. Sriram, "BGP AS\_PATH Verification Based on Autonomous System Provider Authorization (ASPA) Objects",

Work in Progress, Internet-Draft, draft-ietf-sidrops-aspa-verification-24, 19 October 2025,  
 <<https://datatracker.ietf.org/doc/html/draft-ietf-sidrops-aspa-verification-24>>.

[Deprecation-AS\_SET-AS\_CONFED\_SET]

Kumari, W., Sriram, K., Hannachi, L., and J. Haas,  
 "Deprecation of AS\_SET and AS\_CONFED\_SET in BGP", Work in Progress, Internet-Draft, draft-ietf-idr-deprecate-as-set-confed-set-18, 7 March 2025,  
 <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-deprecate-as-set-confed-set-18>>.

[FC-ARXIV] "Secure Inter-domain Routing and Forwarding via Verifiable Forwarding Commitments", September 2023,  
 <<https://arxiv.org/abs/2309.13271>>.

## Appendix A. Attachment

### A.1. Comparison to Other Technologies

#### A.1.1. BGPsec

For basic comparison, please see Section 1.

##### A.1.1.1. Deployment Benefits Analysis

One of the core differences between FC-BGP and BGPsec is the partial deployment scenario. It is difficult for FC-BGP and BGPsec to make an entire deployment, which is an evolution process.

The propagation of the BGP UPDATE message can be simplified as a line. Take the following propagation path as an example.

```

AS(1) --- AS(2) -...- AS(k-1) ---- AS(k) -...- AS(n)
                        \          /
                        \--- AS(m) ---/
  
```

Figure 4: An BGP UPDATE propagation path example.

In this propagation path, the settings are: - A path from AS(1) to AS(n) with N-1 hops, where FC-BGP is deployed consecutively from AS(1) to AS(k-1). AS(k) is the first legacy AS that doesn't enable FC-BGP. - An upgraded AS(n) located after AS(k) tries to validate its received BGP path. - A compromised AS(m) intents to hijack the traffic from AS(n) to AS(1).

Suppose the distance between  $AS(m)$ , i.e., the compromised AS, and  $AS(n)$  is  $L$  hops. The best choice for  $AS(m)$  is to pretend to be a neighbor of  $AS(k)$  and construct a fake path  $AS(1)-AS(2)-\dots-AS(k)-AS(m)-\dots-AS(n)$  with length  $K-1+1+L$  hops.  $AS(m)$  can successfully hijack the traffic only if  $N-1 > K-1+1+L$ , which implies that has to be smaller than  $N-L-1$ .

In the full path deployment scenario, i.e.,  $K=N+1$ , FC-BGP and BGPsec have the same path protection rate. However, in the partial path deployment scenario, i.e.,  $K < N-L-1$ , FC-BGP can protect more paths than BGPsec.

The conclusion can be drawn that FC-BGP provides strictly more security benefits than BGPsec in partial/incremental deployment.

Under full deployment, FC-BGP and BGPsec achieve comparable security benefits as both are capable of preserving the authenticity and immutability of the `AS_PATH` attribute within BGP UPDATE messages. From a security guarantee perspective, FC-BGP establishes a cryptographic commitment and propagates this commitment across the entire network. Through this mechanism, each BGP router can verify that it has received routes from the preceding AS and has forwarded them to the subsequent AS in the path. This approach ensures security through modular pathlets, which are incrementally validated and aggregated to construct a fully authenticated end-to-end routing path. In contrast, BGPsec does not incorporate such a commitment-based framework, relying instead on cryptographic path validation at each AS hop without modular decomposition of path security.

#### A.1.2. ASPA

The ASPA [ASPA-Profile] [ASPA-Verification] mechanism is designed to solve the problem of route leak with RPKI ASPA signed objects and `AS_PATH` attribute. It is an off-path mechanism with a lightweight cryptographic cost for BGP routers. However, it does not protect the `AS_PATH` attribute. Thus, ASPA and FC-BGP are complementary technologies.

#### A.1.3. Only to Customer (OTC) Attribute

OTC is a route-leak detection and prevention mechanism. However, the OTC value itself is not protected. It can be forged. With the signature, FC-BGP can protect the Flags-P2C flag and the Flags-P2P flag. So the FC-BGP route leak prevention mechanism is complementary to the OTC attribute.

## A.2. Implementation Status

We implement the FC-BGP mechanism with FRR version 9.0.1. The implementation includes verifying the FC path attribute upon receiving BGP UPDATE messages and adding and signing the FC path attribute when sending BGP UPDATE messages. The development and testing of this implementation were conducted on Ubuntu 22.04 with OpenSSL 3.X installed.

GitHub repository: <https://github.com/fcbgp/fcbgp-implementation>.

## A.3. An Example

```
AS(65536) --> AS(65537) --> AS(65538)
                        \
                        \--> AS(65539)
```

Figure 5: An FC-BGP UPDATE propagation example.

It is important to note that this example only introduces important steps here and see Section 5.2 for details. What's more, here ASs are global AS and not RS or AS Confederation. No ASPP is taken into consideration.

For the sake of discussion, we assume that AS 65537 receives an FC-BGP UPDATE message for prefix 192.0.2.0/24 from AS 65536 and will send the route to AS 65538 and AS 65539 as Figure 5 shows. An FC-BGP speaker SHOULD propagate an FC-BGP UPDATE message to downstream ASs only after completing the validation and best route path selection.

When the FC-BGP speaker in AS 65536 plans to propagate routes to its downstream AS 65537, it fills the FC path attribute first. The PASN is 0/NULL, the CASN is 65536, and the NASN is 65537. Then it calculates the signature, fills the FC, and forms the FC path attribute and FC-BGP UPDATE message. Then, it sends the UPDATE message out.

When receiving an UPDATE message from AS 65536, the FC-BGP speaker in AS 65537 retrieves the FC path attribute and extracts the FC list. It then finds the FC with CASN == 65536 as the AS\_PATH has only one AS(65536) and checks whether PASN is 0 as well as NASN is 65537. If so, it uses the SKI field to find the public key and calculates the signature using the algorithm specified in the Algorithm ID. If the calculated signature matches the signature in the FC segment, then the AS-Path hop associated with the AS 65536 is verified. This process repeats for all FCs and AS-Paths in the FC list if it has other ASs in AS\_PATH. However, if AS 65537 does not support FC-BGP, the BGP speaker of AS 65537 simply forwards the BGP UPDATE to its neighbors when propagating this FC-BGP route without validating the FC path attribute.

FC-BGP speakers need to generate different UPDATE messages for different neighbors. Each UPDATE announcement contains only one route prefix and cannot be aggregated. This is because different route prefixes may have different announcement paths due to different routing policies. Multiple aggregated route prefixes may cause FC generation and verification errors. When multiple route prefixes need to be announced, the FC-BGP speaker needs to generate different UPDATE messages for each route prefix. Thus, the FC-BGP speaker of AS 65537 generates different UPDATE messages for AS 65538 and AS 65539 separately. The biggest difference is that the NASN is 65538 for AS 65538 and 65539 for AS 65539 in the FC segment generated by AS 65537.

Take AS 65538 as the next hop. The FC-BGP speaker in AS 65537 will encapsulate each prefix to be sent to AS 65538 in a single UPDATE message, add the FC path attribute, and sign the path content using its private key to fill a new FC segment. The FC path attribute and the FC segment use the message format shown in Figure 1 and Figure 2 separately. When signing, the FC-BGP speaker computes the SHA256 hash in the order of (PASN ( 0 if absent), CASN, NASN, IP Prefix Address, and IP Prefix Length) first. Next, the FC-BGP speaker should calculate the digest information Digest, sign the Digest with ECDSA, and then fill in the Signature field and other FC fields. After that, AS 65537 prepends its own FC on top of the FC List. At this point, the processing of FC path attributes by the FC-BGP speaker is complete. The subsequent processing of BGP messages follows the standard BGP process.

## Acknowledgments

The authors would like to thank Keyur Patel, Jeffery Hass, Randy Bush, Maria Matejka, Tobias Fiebig, Nan Geng, Tom Strickx, Susan Hares, R端diger Volk, Jun Zhang, Kotikalapudi Sriram, John Scudder, Job Snijders, Russ Housley, and Andrew for their review and valuable comments.

## Authors' Addresses

Ke Xu  
Tsinghua University  
Beijing  
China  
Email: xuke@tsinghua.edu.cn

Xiaoliang Wang  
Tsinghua University  
Beijing  
China  
Email: wangxiaoliang0623@foxmail.com

Zhuotao Liu  
Tsinghua University  
Beijing  
China  
Email: zhuotaoliu@tsinghua.edu.cn

Qi Li  
Tsinghua University  
Beijing  
China  
Email: qli01@tsinghua.edu.cn

Jianping Wu  
Tsinghua University  
Beijing  
China  
Email: jianping@cernet.edu.cn

Yangfei Guo  
Zhongguancun Laboratory  
Beijing  
China  
Email: [guoyangfei@zgclab.edu.cn](mailto:guoyangfei@zgclab.edu.cn)