

Routing Area Working Group
Internet-Draft
Intended status: Standards Track
Expires: 8 January 2026

Z. Ma
Y. Li
CNIC, CAS
L. Lu
China Mobile
G. Xie
CNIC, CAS
7 July 2025

BGP Extensions for Service Routing of Mobile Core Networks
draft-ma-bgp-mobile-core-extensions-00

Abstract

This document describes a new route propagation and service discovery mechanism by proposing BGP extensions for mobile core network service routing. The proposed solution introduces a Service Route Address Family, hierarchical Network Identifier allocation, and path-aware routing mechanisms that enable scalable inter-domain service discovery while preserving compatibility with current 3GPP standards. These extensions provide the foundation for efficient service propagation, route aggregation, and loop-free routing in large-scale distributed mobile core deployments.

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 8 January 2026.

Copyright Notice

Copyright (c) 2025 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
2. Terminology	3
3. Problem Statement	4
3.1. Current Control Plane Limitations	4
3.1.1. HTTP-Based Service Discovery and Propagation Limitations	4
3.1.2. NID Allocation and Management	5
3.2. Example Network Topology Illustrating Current Limitations	6
4. BGP Extensions	8
4.1. Hierarchical Network Identifier (NID) Extension	8
4.1.1. Assignment Mode 3 Extension	8
4.1.2. Hierarchical Encoding Scheme	8
4.1.3. NID Allocation Example	9
4.2. NID Attribute	9
4.2.1. Attribute Definition	10
4.2.2. Encoding Format	10
4.2.3. Processing Rules	10
4.3. NID_PATH Attribute	11
4.3.1. Attribute Definition	11
4.3.2. Encoding Format	11
4.3.3. Processing Rules	11
4.4. Service Route AFI/SAFI	12
4.4.1. AFI/SAFI Definition	12
4.4.2. Service Route NLRI	12
4.5. Domain Routes	15
4.5.1. Definition and Purpose	15
4.5.2. Use Case and Motivation	15
4.5.3. Domain Route Characteristics	15
4.5.4. Domain Route Format	16
4.6. Operation	16
4.6.1. BGP Session Establishment	16
4.6.2. Service Information Base (SIB)	17
4.6.3. Generation of Service Routes	18
4.6.4. Service Route Transmission	19
4.6.5. Service Route Withdrawal	20
4.6.6. Partial Service Route Updates	21

5. Security Considerations	23
5.1. BGPsec Integration	23
6. IANA Considerations	23
6.1. BGP AFI/SAFI Assignment	23
6.2. BGP Path Attribute Type Codes	24
6.3. Service Route ID Format	24
7. References	24
7.1. Normative References	24
7.2. Informative References	25
Authors' Addresses	25

1. Introduction

Service discovery in the 5G core network control plane is primarily communicated through the HTTP/2 protocol [TS29.500]. The 3GPP Service-Based Architecture (SBA) defined in [TS23.501] enables flexible deployment models, including multi-access edge computing (MEC), hybrid cloud scenarios, and distributed network function placements. However, these distributed deployments fundamentally challenge existing inter-domain communication and service discovery mechanisms. Current HTTP-based control plane protocols were designed for centralized environments and lack the path awareness, aggregation capabilities, and scalability required for distributed operations.

This document proposes a new service discovery mechanism by proposing BGP extensions for mobile core network service routing. The specification introduces a Service Route Address Family Identifier (AFI) and Subsequent Address Family Identifier (SAFI), hierarchical Network Identifier (NID) allocation schemes, and the NID_PATH attribute for path-aware routing and loop prevention. These extensions transform BGP into a capable control plane protocol for mobile core networks, enabling efficient inter-domain service discovery, intelligent route aggregation, and reliable propagation of network function capabilities across distributed domains.

2. Terminology

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.

This document uses the following terms:

***Network Identifier (NID)*:** A hierarchical identifier used to uniquely identify network domains, extending the concept defined in [TS23.501].

- *NID_PATH*: A BGP path attribute that records the sequence of NIDs traversed by a route in non-BGPsec environments, providing basic loop prevention and path traceability. This attribute is superseded by BGPsec's BGPsec_PATH when cryptographic protection is available.
- *Service Route*: A new BGP Address Family Identifier (AFI) and Subsequent Address Family Identifier (SAFI) for propagating mobile core network service information. Service routes carry aggregated service-level information derived from multiple NF instances.
- *Service Route ID*: A 16-byte UUID version 7 identifier that uniquely identifies a complete service route across all its fragments.
- *Service Route Fragment*: A portion of a service route that fits within a single BGP UPDATE message. When a service route is too large for one UPDATE message, it is split into multiple fragments.
- *Service Attributes*: Network Function attributes that describe the properties and capabilities of services offered by a network domain, such as PLMN, SST, SD, locality, etc.
- *User Equipment (UE)*: A device that provides access to network services for a human user or an application, as defined in [TS23.501].

3. Problem Statement

The evolution towards large-scale distributed 5G mobile core networks introduces fundamental challenges that existing HTTP-based control plane mechanisms cannot adequately address. These challenges span from basic connectivity and discovery to advanced requirements for multi-path optimization, load balancing, and security.

3.1. Current Control Plane Limitations

3.1.1. HTTP-Based Service Discovery and Propagation Limitations

Current 5G core networks rely on HTTP-based route propagation and service discovery mechanisms through the Network Repository Function (NRF) as defined in [TS29.510]. While this approach works adequately in centralized deployments, it presents fundamental architectural limitations that become critical in large-scale distributed environments.

HTTP-based route propagation provides no path or topology information, making it impossible to optimize routing decisions or understand network connectivity patterns. Without this fundamental routing information, networks cannot make informed decisions about service placement or access paths. This topology blindness forces networks to rely on static configuration or sub-optimal fallback mechanisms when multiple service paths are available.

Each Network Function (NF) profile is propagated individually without aggregation capabilities, leading to excessive signaling overhead and inefficient route lookup operations. This lack of aggregation becomes particularly problematic as networks scale, creating bandwidth consumption issues and processing bottlenecks at NRF entities.

3.1.1.1. BGP as a Solution Foundation

The Border Gateway Protocol (BGP) directly addresses these HTTP-based limitations through proven mechanisms that have been refined over decades of Internet operation. BGP's path vector architecture provides complete topology information and loop prevention capabilities that are essential for distributed networks, directly addressing the fundamental topology awareness gaps in current HTTP-based systems.

BGP includes native support for route aggregation, enabling efficient summarization of service information and reducing signaling overhead. The protocol's proven scalability to global Internet routing, makes it well-suited for large-scale distributed core networks that may encompass thousands of domains and millions of service endpoints.

Additionally, BGP benefits from a rich ecosystem of security extensions through BGPsec, multi-path routing capabilities, traffic engineering mechanisms, and sophisticated policy enforcement frameworks. Rather than attempting to retrofit the existing HTTP-based control plane with complex modifications to address these fundamental limitations, adopting BGP directly provides a proven, immediately available solution that inherits decades of Internet-scale operational experience and continuous enhancement.

3.1.2. NID Allocation and Management

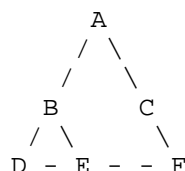
The existing NID definition in [TS23.501] uses a flat 44-bit identifier space with significant limitations that hinder scalable network deployment. NID allocation requires coordination between operators or assignment by IANA, creating operational overhead and deployment delays that prevent dynamic network expansion. This manual coordination requirement becomes particularly problematic in

cloud-native deployments where network domains may need to be created and destroyed dynamically in response to changing traffic patterns or infrastructure requirements.

The flat NID structure lacks semantic information about network hierarchy or domain relationships, making network management, troubleshooting, and automated route aggregation impossible. Without hierarchical semantics, network operators cannot implement efficient routing policies or perform automatic aggregation based on organizational or geographical boundaries. Additionally, static allocation schemes cannot support dynamic domain creation and deletion required for modern cloud-native deployments and edge computing scenarios, where network domains may need to scale elastically based on demand.

3.2. Example Network Topology Illustrating Current Limitations

To illustrate these limitations, consider the following hierarchical network topology with both Parent-to-Child (P2C) and Peer-to-Peer (P2P) relationships:



Network Topology:

- * Node A: Central domain (root)
- * Nodes B, C: Regional domains (children of A, peers to each other)
- * Nodes D, E, F: Edge domains (D and E are children of B, F is child of C)
- * Direct peer connection exists between E and F

Routing Table of Each Node Example:

Node	Next Hop	Route Info
A	B	B&D&E
A	C	C&F
B	D	D
B	E	E
C	F	F
D	E	E
E	D	D
F	E	E

Table 1

Current HTTP-Based System Limitations:

The current HTTP-based system exhibits severe limitations in topology-aware path selection. When domain A needs to consume services, both domains B and F can provide the required service. However, without topology information, A cannot determine that B is topologically closer (1 hop) compared to F (2 hops via C), leading to suboptimal service selection that wastes network resources and increases latency.

Loop prevention failures in peer-to-peer propagation represent another critical weakness. When domain D needs to discover services from domain F, the lack of loop prevention mechanisms prevents route information from being safely propagated along P2P paths. Without an AS_PATH equivalent, route information cannot be transmitted from F through the E-F peer connection to E, then to D, forcing D to rely on the inefficient hierarchical path D -> B -> A -> C -> F instead of the optimal direct peer path D -> E -> F.

The absence of route aggregation capabilities creates both scaling and security concerns. Parent domain A receives individual, non-aggregated NF profiles from all downstream domains (B, C, D, E, F), exposing detailed internal topology and creating excessive routing overhead. This lack of aggregation not only increases bandwidth consumption but also reveals sensitive network topology information to upstream domains.

Security vulnerabilities in route propagation present significant operational risks. Each NRF can freely modify and tamper with route information received from peers before propagating it to other peers. Without cryptographic protection mechanisms, there is no guarantee of route integrity or authenticity, allowing malicious or compromised NRFs to inject false service information or manipulate routing decisions.

Finally, the missing path metrics capability prevents intelligent routing decisions. HTTP-based discovery provides no path cost, latency, or reliability information that would enable intelligent path selection based on network conditions, forcing networks to make routing decisions without crucial performance data.

4. BGP Extensions

4.1. Hierarchical Network Identifier (NID) Extension

4.1.1. Assignment Mode 3 Extension

This document extends assignment mode 3 of the NID specification to support hierarchical allocation. The hierarchical structure enables efficient route aggregation at parent domains while allowing automatic sub-domain creation without external coordination. This approach provides clear representation of domain relationships that can be leveraged for routing optimization and network management purposes.

4.1.2. Hierarchical Encoding Scheme

The NID hierarchy supports up to three levels in the current specification, with provisions for extension to support additional levels through the reserved field. The reserved field functions as a hierarchy level indicator, enabling the scheme to scale beyond three levels when needed for complex network deployments. This flexible approach allows networks to adopt appropriate hierarchy depths based on their organizational structure and operational requirements.

Level	Mode Bits	Reserved	Address Space Distribution
1	3 (011)	1 (0)	39 bits for global domains
2	3 (011)	2 (10)	12 bits + 26 bits for regional/national domains
3	3 (011)	3 (110)	7 bits + 12 bits + 18 bits for local/edge domains
4+	3 (011)	4+ (1110+)	Future extension for deeper hierarchies

Table 2

The reserved field serves as a hierarchy level indicator, enabling the scheme to scale beyond three levels when needed for complex network deployments.

4.1.3. NID Allocation Example

Consider NID 3.6.1.1.0 encoded as:

```
0011 110 0000001 000000000001 00000000000000000001
```

This represents:

- * Assignment mode: 3 (0011)
- * Hierarchy level: 3 (110)
- * Level 1 domain: 1 (0000001)
- * Level 2 domain: 1 (000000000001)
- * Level 3 domain: 1 (00000000000000000001)

The parent domain is 3.6.1.0.0, and this domain can allocate sub-domains following the pattern 3.6.1.1.x.

4.2. NID Attribute

4.2.1. Attribute Definition

The NID attribute is an optional transitive BGP path attribute that carries the Network Identifier of the originating domain. This attribute enables receiving domains to identify the source of service routes and make routing decisions based on domain hierarchy and relationships.

4.2.2. Encoding Format

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
| Attr. Flags |Attr. Type Code|           Attr. Length           |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     NID (6 octets)                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                                                                                                                 |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

* *Attribute Flags*: Optional (1), Transitive (1), Partial (0), Extended Length (0)

* *Attribute Type Code*: TBD (to be assigned by IANA)

* *Attribute Length*: 6 (fixed length for NID)

* *NID*: 6-byte Network Identifier of the originating domain

4.2.3. Processing Rules

When processing a route with NID attribute, receiving domains must follow specific procedures to maintain route integrity and enable proper policy enforcement. The NID attribute identifies the originating domain of the service route, allowing receiving domains to make routing decisions based on the source domain's identity and characteristics. Receiving domains may apply import and export policies based on the originating NID, enabling fine-grained control over which service routes are accepted and propagated.

The receiving domain may validate the NID against known domain hierarchy for operational purposes, though such validation is not required for basic protocol operation. This validation can help detect configuration errors or potential security issues. The NID attribute may be used in best path selection algorithms to prefer routes from specific domains or hierarchy levels, allowing networks to implement policies that favor certain service providers or geographic regions based on operational requirements.

4.3. NID_PATH Attribute

4.3.1. Attribute Definition

The NID_PATH attribute is an optional non-transitive BGP path attribute that provides basic loop prevention and path traceability in environments where BGPsec is not deployed. In BGPsec-enabled networks, this attribute is not necessary as BGPsec's BGPsec_PATH attributes provide superior cryptographic protection for path validation.

The NID_PATH attribute serves two primary functions in non-BGPsec environments:

1. ***Loop Prevention***: Detects and prevents routing loops by checking if a domain's NID already appears in the path.
2. ***Path Traceability***: Maintains complete propagation history for troubleshooting and audit purposes.

4.3.2. Encoding Format

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9
Attr. Flags										Attr. Type Code										Attr. Length																			
NID List (variable)																																							

* ***Attribute Flags***: Optional (0), Non-transitive (0), Complete (1), Extended Length (as needed)

* ***Attribute Type Code***: TBD (to be assigned by IANA)

* ***Attribute Length***: Variable, depending on the number of NIDs

* ***NID List***: Sequence of 6-byte NIDs in reverse propagation order

4.3.3. Processing Rules

When processing a route with NID_PATH, receiving domains must implement loop detection and path management procedures to ensure routing stability. If the receiving domain's NID appears in the path, the route must be discarded to prevent routing loops that could destabilize the network. This loop detection mechanism is critical for maintaining network stability, particularly in complex topologies with multiple interconnection points between domains.

When propagating the route to other peers, the domain must prepend its own NID to the path before transmission. This prepending operation maintains the complete propagation history and enables downstream domains to perform their own loop detection. Receiving domains may validate the path against known topology for operational purposes, though such validation is optional and intended primarily for troubleshooting and audit functions.

4.4. Service Route AFI/SAFI

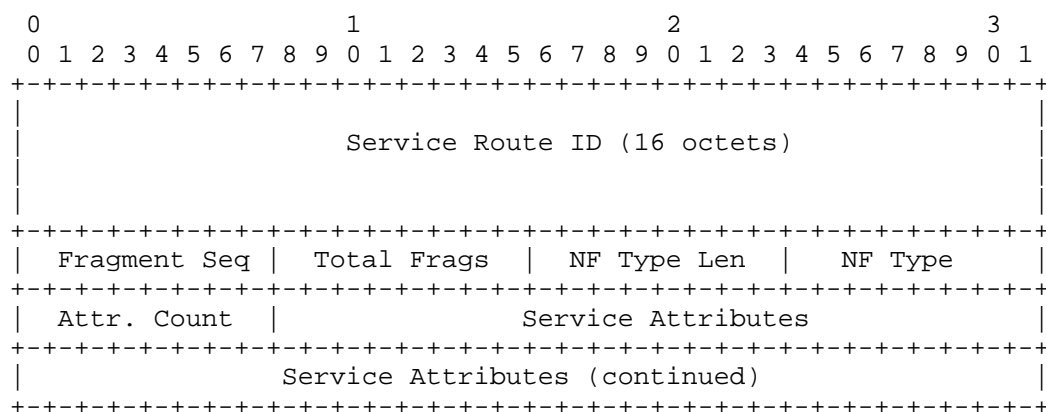
4.4.1. AFI/SAFI Definition

A new AFI (value TBD) and SAFI (value TBD) are defined for Service Route with the following characteristics:

- * *Address Family*: Service Route (TBD)
- * *SAFI Value*: TBD

4.4.2. Service Route NLRI

Service Routes carry aggregated service-level information through Service Route Fragments. Due to the large size of core network routing entries, a fragmentation mechanism is used to split service routes across multiple UPDATE messages.



4.4.2.1. Service Route ID

A 16-byte UUID version 7 identifier that uniquely identifies a complete service route. All fragments belonging to the same service route share the same Service Route ID.

4.4.2.2. Service Route Fragment

When a service route is too large to fit in a single UPDATE message, it is split into multiple Service Route Fragments:

- * *Fragment Seq*: The sequence number of this fragment (1-based)
- * *Total Frags*: The total number of fragments for this service route

4.4.2.3. NF Type

A single Network Function type supported by this fragment, as defined in [TS23.501]:

- * AMF (Access and Mobility Management Function)
- * SMF (Session Management Function)
- * UPF (User Plane Function)
- * PCF (Policy Control Function)
- * UDM (Unified Data Management)
- * UDR (Unified Data Repository)
- * NRF (Network Repository Function)
- * NSSF (Network Slice Selection Function)
- * AUSF (Authentication Server Function)
- * etc.

The complete list of Network Function types is specified in [TS23.501].

4.4.2.4. Service Attributes

Each Service Attribute follows this structure:

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								
-----Attribute Type (4 octets)-----																																							
Flags										Defined By Length										Defined By																			
Value Len										Values																													

***Attribute Types*:**

The following attribute types are mapped to numeric values for encoding in the Attribute Type field, based on Network Function Profile attributes defined in [TS29.510]:

1. PLMN (Public Land Mobile Network)
2. SST (Slice/Service Type)
3. SD (Slice Differentiator)
4. Locality
5. GPSI (Generic Public Subscription Identifier)
6. SUPI (Subscription Permanent Identifier)
7. Routing Indicator
8. TAC (Tracking Area Code)
9. DNN (Data Network Name)

Additional attribute types SHALL be defined to cover all NF Profile attributes specified in [TS29.510] and related 3GPP standards.

***Flags* :**

- * Optional (O): Whether this attribute is optional for matching
- * Transitive (T): Whether this attribute should be propagated
- * Encode Type (ET): 3-bit field indicating value encoding format
 - 000: Value is a numeric list
 - 001: Value is a string list

- 010-111: Reserved for future use

4.5. Domain Routes

4.5.1. Definition and Purpose

Domain Routes represent a special category of service routes designed to address specific privacy and accessibility requirements in distributed mobile core networks. A Domain Route contains only PLMN (Public Land Mobile Network) information as its service attribute and uses a fixed all-zero Service Route ID (00000000-0000-0000-0000-000000000000).

4.5.2. Use Case and Motivation

Domain Routes serve critical scenarios where private network domains require selective visibility within the broader network ecosystem. In certain deployments, private network domains prefer to maintain operational privacy by not exposing detailed service capabilities or internal network function profiles to external domains. However, these private domains still need to enable User Equipment (UE) that belongs to their PLMN to access services when roaming or connecting from external network locations.

The primary use case involves UEs whose subscription services or information are associated with a private network domain but are currently registered to a public network domain. The public network needs to discover and access subscription services for these UEs within their home private domain through NID and PLMN. Without Domain Routes, external domains would have no knowledge of the private domain's existence or basic accessibility information, preventing proper service routing for roaming UEs.

4.5.3. Domain Route Characteristics

Domain Routes have several distinctive characteristics that differentiate them from regular service routes:

***Minimal Information Disclosure*:** Domain Routes contain only essential PLMN information, revealing no details about internal network functions, service capabilities, or network topology. This approach preserves privacy while enabling basic connectivity.

***Fixed Route Identifier*:** All Domain Routes use the same all-zero Service Route ID (00000000-0000-0000-0000-000000000000), simplifying identification and processing across network domains.

***No Fragmentation*:** Due to their minimal content, Domain Routes never require fragmentation and always fit within a single BGP UPDATE message.

***Universal Propagation*:** Domain Routes are typically propagated more broadly than detailed service routes, ensuring that UE accessibility information reaches all relevant network domains.

4.5.4. Domain Route Format

A Domain Route follows the standard Service Route NLRI format but with specific field values:

0																1																2																3															
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9																								
All-Zero Service Route ID (16 octets)																																																															
00000000-0000-0000-0000-000000000000																																																															
Fragment=1																Total=1																NF Type Len																NF Type															
Attr.Count=1																PLMN Attribute																																															
PLMN Attribute (continued)																																																															

***Field Values for Domain Routes*:**

- * ***Service Route ID*:** Fixed value
00000000-0000-0000-0000-000000000000
- * ***Fragment Seq*:** Always 1 (single fragment)
- * ***Total Frags*:** Always 1 (no fragmentation)
- * ***NF Type*:** Set to a reserved value indicating "Domain Route" (TBD)
- * ***Attr. Count*:** Always 1 (only PLMN attribute)
- * ***Service Attributes*:** Contains only one PLMN attribute with the domain's PLMN information

4.6. Operation

4.6.1. BGP Session Establishment

4.6.1.1. Peer Establishment Between NRF Entities

When two NRF entities need to establish a BGP session for service route exchange, they follow the standard BGP session establishment procedure with extensions specific to mobile core network requirements. Instead of using traditional Autonomous System (AS) numbers, NRF entities use their respective NIDs for session identification and loop prevention, aligning the session establishment with the hierarchical domain structure used throughout the mobile core network.

The BGP OPEN message must include several mandatory elements to support Service Route operations. The local domain's NID is included in place of the AS number field, establishing the domain identity for the session. Capability advertisement for Service Route AFI/SAFI support is required to indicate that the peer can process Service Route messages.

The OPEN message format for NRF BGP sessions:

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								
Version										My NID (first 2 octets)																													
										My NID (remaining 4 octets)																													
Hold Time										BGP Identifier																													
										BGP Identifier (continued)																													
Opt Parm Len										Optional Parameters																													

4.6.2. Service Information Base (SIB)

4.6.2.1. SIB Structure and Management

Each NRF entity maintains a Service Information Base (SIB) that stores service route information in an organized, efficient manner. The SIB contains local service routes originated by the local domain based on registered NF instances, remote service routes learned from BGP peers, and route state information including metadata such as route origin, next hop, and propagation path. This comprehensive information storage enables efficient route lookup and policy application during service discovery operations.

4.6.2.2. SIB Entry Format

Each SIB entry must contain the following essential information elements:

- * Service Route ID for unique identification across fragments
- * Origin NID identifying the domain that originally advertised the route
- * Next Hop NID for forwarding decisions
- * NID_PATH for loop prevention and path traceability (only in non-BGPsec environments)
- * Fragment information including total fragment count and completion status
- * Network Function type and associated service attributes
- * Route status and timestamp information for operational management

This information enables efficient route lookup, policy application, and fragment reassembly during service discovery operations. In BGPsec-enabled deployments, cryptographic path information is maintained through BGPsec's BGPsec_PATH attributes instead of the NID_PATH field.

4.6.3. Generation of Service Routes

4.6.3.1. NF Registration and Aggregation Process

When NFs register with the local NRF, a systematic process generates service routes that aggregate service capabilities across multiple NF instances. The NRF collects NF profiles from registered NF instances according to [TS29.510], ensuring compliance with existing 3GPP standards for NF registration and profile management.

For each NF type, the NRF aggregates all unique service attributes across registered instances of that type. This aggregation process creates a comprehensive view of the services available within the domain for each Network Function type. The aggregation algorithm examines all NF instances of a given type and collects their unique service attributes, avoiding duplication while ensuring complete coverage of available capabilities.

Each aggregated service becomes a single service route with a unique Service Route ID generated using UUID version 7. The generated service route is stored in the local SIB and marked for propagation to appropriate BGP peers based on configured export policies.

4.6.3.2. Service Route Fragmentation

If a service route exceeds the maximum BGP UPDATE message size, it must be fragmented using a systematic approach that maintains route integrity while enabling efficient transmission. The service route is split into multiple fragments, where each fragment contains the same Service Route ID, sequential fragment numbers, total fragment count, and a subset of service attributes. This fragmentation approach ensures that receiving domains can properly reassemble the complete service route while processing fragments in any order.

Fragmentation follows specific rules to maintain consistency and enable efficient processing. The minimum fragment unit consists of one NF Type plus one Service Attribute, ensuring that each fragment contains meaningful information. Each fragment contains exactly one NF Type to simplify processing logic, while Service Attributes can be distributed across fragments as needed to stay within message size limits. All fragments share the same Service Route ID, enabling receiving domains to associate fragments with the correct service route during reassembly.

4.6.4. Service Route Transmission

4.6.4.1. MP_REACH_NLRI Propagation

Service routes are propagated using MP_REACH_NLRI through a systematic procedure that ensures efficient and reliable route distribution. The NRF selects service routes from the SIB for propagation based on multiple criteria including route origin (local vs. learned), configured export policies, and peer relationships. This selection process ensures that only appropriate routes are propagated to each peer according to established network policies.

For each selected route, the NRF performs attribute processing to maintain routing integrity and enable proper loop prevention. The system adds or updates the NID attribute with the local domain NID, identifying this domain as the most recent propagator of the route. The local NID is prepended to the NID_PATH attribute to maintain the complete propagation history. Any configured attribute modifications are applied according to local policy requirements.

When fragmentation is required, the system generates multiple UPDATE messages with specific handling characteristics. Each UPDATE contains one fragment of the original service route, and fragments may be transmitted in any order since no inter-fragment dependencies exist. This approach enables efficient parallel transmission and processing of large service routes while maintaining protocol simplicity.

4.6.5. Service Route Withdrawal

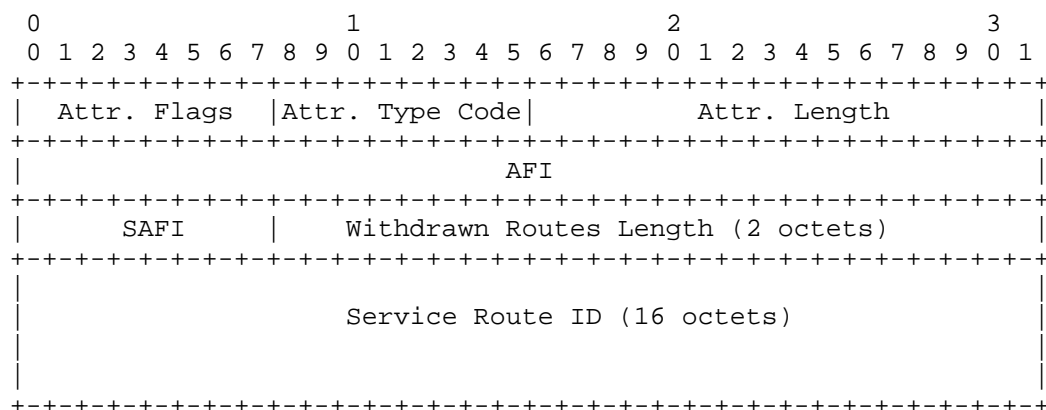
4.6.5.1. Withdrawal Procedure

To withdraw a complete service route, the originating NRF sends an MP_UNREACH_NLRI containing only the Service Route ID, providing a simple and efficient withdrawal mechanism. The MP_UNREACH_NLRI contains the appropriate AFI/SAFI for Service Route and the NLRI consisting of only the Service Route ID (16 octets). This streamlined approach eliminates the need to transmit detailed route information during withdrawal operations.

Receiving NRFs process withdrawals by locating all fragments with matching Service Route ID and removing them from the SIB. The withdrawal implicitly affects all fragments belonging to the service route, regardless of how many fragments were originally received. The receiving NRF then propagates the withdrawal to appropriate peers according to configured export policies, ensuring that the withdrawal information reaches all domains that previously received the service route.

4.6.5.2. Withdrawal Message Format

The MP_UNREACH_NLRI attribute for service route withdrawal has the following detailed format:



Field Descriptions:

- * ***Attribute Flags (8 bits)*:** Optional (1), Transitive (0), Partial (0), Extended Length (0)
- * ***Attribute Type Code (8 bits)*:** MP_UNREACH_NLRI (15)
- * ***Attribute Length (16 bits)*:** Total length of the attribute value (23 octets)
- * ***AFI (16 bits)*:** Address Family Identifier for Service Route (TBD)
- * ***SAFI (8 bits)*:** Subsequent Address Family Identifier for Service Route (TBD)
- * ***Withdrawn Routes Length (16 bits)*:** Length of withdrawn routes field (16 octets)
- * ***Service Route ID (128 bits)*:** 16-byte UUID v7 identifier of the service route to withdraw

Processing Notes:

The withdrawn routes field contains only the Service Route ID, providing a simple and unambiguous identification of the route to be withdrawn. No additional NLRI information is required for withdrawal since the Service Route ID uniquely identifies the complete service route across all its fragments. All fragments sharing the same Service Route ID are implicitly withdrawn when this withdrawal message is processed, ensuring complete cleanup of the service route from the receiving domain's routing tables.

4.6.6. Partial Service Route Updates

4.6.6.1. Incremental Update Mechanism

For efficiency, NRFs can perform partial updates of service routes without withdrawing and re-advertising the entire route, reducing network overhead and improving convergence times. When service attributes change, the system generates new fragments with the same Service Route ID, includes only modified attributes in the update, and increments fragment sequence numbers if needed to maintain proper ordering.

The attribute replacement mechanism follows specific semantics to ensure consistency across all receiving domains. The update contains the same Service Route ID as the original route, updated service attributes, and follows replacement semantics rather than additive

semantics. This approach ensures that receiving domains can precisely understand which attributes have changed and how to update their local routing tables.

Receiving NRFs process partial updates by identifying the existing route by Service Route ID, replacing specified attributes entirely while maintaining unchanged attributes from previous advertisements, and propagating updates according to local policies. This processing approach ensures that partial updates are applied consistently across the network while preserving routing table integrity.

4.6.6.2. Update Message Semantics

Partial updates follow specific rules that ensure consistent interpretation across all network domains while maintaining routing table integrity. Updates operate at the service attribute level, providing fine-grained control over which aspects of a service route are modified. Updated attributes completely replace previous values rather than being merged with existing values, eliminating ambiguity about the final state of modified attributes. Non-updated attributes remain unchanged during partial update operations, ensuring that unrelated service characteristics are preserved.

4.6.6.3. Example Update Sequence

Consider a service route update scenario where an AMF service needs to support an additional network slice:

Initial Advertisement:

Service Route ID: 550e8400-e29b-41d4-a716-446655440000
NF Type: AMF
Attributes: PLMN=001-01, SST=1,2, Locality=Region1

Partial Update:

Service Route ID: 550e8400-e29b-41d4-a716-446655440000
NF Type: AMF
Attributes: SST=1,2,3

Resulting State:

Service Route ID: 550e8400-e29b-41d4-a716-446655440000
NF Type: AMF
Attributes: PLMN=001-01, SST=1,2,3, Locality=Region1

In this example, only the SST attribute is updated to include slice type 3, while the PLMN and Locality attributes remain unchanged from the original advertisement.

5. Security Considerations

The BGP extensions proposed in this document inherit the security framework of the base BGP protocol and can leverage existing BGP security mechanisms to ensure route integrity and authenticity.

5.1. BGPsec Integration

The Service Route extensions defined in this document are designed to be compatible with BGPsec, which provides cryptographic protection for BGP route advertisements. BGPsec can be applied to Service Route advertisements to ensure:

***Route Origin Authentication*:** BGPsec validates that the originating NRF domain is authorized to advertise specific service routes, preventing unauthorized route injection. Each Service Route advertisement includes cryptographic signatures that verify the identity of the originating domain using its NID.

***Path Validation*:** BGPsec's BGPsec_PATH attributes provide cryptographic path validation that supersedes the need for the NID_PATH attribute. When BGPsec is deployed, the BGPsec_PATH mechanism ensures that each domain in the propagation path has legitimately forwarded the route through cryptographic signatures, providing both loop prevention and path authentication. In BGPsec deployments, the NID_PATH attribute becomes redundant as the BGPsec_PATH provides superior cryptographic protection for path validation and traceability.

***Route Integrity*:** BGPsec protects Service Route content from modification during propagation. Service attributes, fragmentation information, and route identifiers are cryptographically protected, ensuring that receiving domains can trust the accuracy of service information.

***Non-repudiation*:** BGPsec provides non-repudiation capabilities that enable audit trails for service route propagation. Network operators can verify the complete history of route advertisements and identify the source of any security incidents.

6. IANA Considerations

6.1. BGP AFI/SAFI Assignment

IANA is requested to assign a new AFI value for Service Route from the "Address Family Numbers" registry.

* Suggested value: TBD

- * Reference: This document

IANA is requested to assign a new SAFI value for Service Route from the "Subsequent Address Family Identifiers (SAFI) Parameters" registry.

6.2. BGP Path Attribute Type Codes

IANA is requested to assign new path attribute type codes from the "BGP Path Attributes" registry:

1. NID attribute for network domain identification

- * Suggested value: TBD
- * Reference: This document

2. NID_PATH attribute for loop prevention and path traceability

- * Suggested value: TBD
- * Reference: This document

6.3. Service Route ID Format

This document uses UUID version 7 as defined in RFC 9562 for Service Route ID generation. No IANA action is required for UUID usage.

7. References

7.1. Normative References

- [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>>.
- [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>>.
- [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>>.

- [RFC9562] Davis, K., Peabody, B., and P. Leach, "Universally Unique IDentifiers (UUIDs)", RFC 9562, DOI 10.17487/RFC9562, May 2024, <<https://www.rfc-editor.org/rfc/rfc9562>>.

7.2. Informative References

- [TS23.501] 3GPP, "System architecture for the 5G System (5GS)", 3GPP TS 23.501 , 2023.
- [TS29.500] 3GPP, "5G System; Technical Realization of Service Based Architecture", 3GPP TS 29.500 , 2022.
- [TS29.510] 3GPP, "5G System; Network Function Repository Services", 3GPP TS 29.510 , 2023.

Authors' Addresses

Zhuoran Ma
CNIC, CAS
Email: mazhuoran@hnu.edu.cn

Yanbiao Li
CNIC, CAS
Email: lybmath@cnic.cn

Lu Lu
China Mobile
Email: lulu@chinamobile.com

Gaogang Xie
CNIC, CAS
Email: xie@cnic.cn