6man
Internet-Draft
Intended status: Standards Track
Expires: 11 June 2026

X. Gu, Ed.
N. Zhang, Ed.
X. Yi, Ed.
China Unicom

8 December 2025

A whitelist-based data fencing mechanism in data circulation
draft-gu-6man-whitelist-data-circulation-00

Abstract

This document defines the data transaction permission fields within the Data Fencing architecture and the whitelist information maintained by the gateway devices. By comparing the data transaction permission fields against the whitelist, data packets are either permitted or blocked, establishing precise data flow boundaries and egress control. Finally, the document presents several use cases of the data fencing.

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 11 June 2026.

Copyright Notice

Table of Contents

1. Introduction

Amidst the accelerating wave of digital transformation, data has undeniably emerged as a core factor of production, whose value can only be maximized through sufficient circulation, sharing, and integration. However, this very flow and sharing of data introduce formidable challenges to security and privacy preservation. Traditional data access control mechanisms, such as firewalls, Virtual Private Networks (VPNs), and Role-Based Access Control (RBAC), have played significant roles in static, well-defined network environments. Yet, in complex scenarios like cross-domain data exchange, multi-cloud collaboration, and industrial chain coordination, data becomes susceptible to replication, download, and forwarding. Once data traverses traditional security boundaries, these static policies prove inadequate, failing to meet the demands of dynamic data flow. This inadequacy precipitates a sharp increase

in risks, including data misuse, unauthorized dissemination, and insider threats.

In this context, Data Fencing has emerged as a novel paradigm for data security. It moves beyond reliance on fixed network perimeters by deeply integrating security policies with the data itself, thereby establishing a dynamic, mandatory security boundary that persists throughout the entire data lifecycle. The core principle of Data Fencing lies in employing policy definitions and technical measures to enforce granular control and restrictions over data access, transmission, and usage scope. This ensures that sensitive data remains protected from access or leakage by unauthorized users or systems.

Within this architecture, the gateway device assumes a critical role as the policy enforcement point. By dynamically maintaining whitelists, it constructs granular boundaries for data flow. This mechanism ensures that data moves securely within authorized confines, effectively guarding against unauthorized access and insider threats. Consequently, it upholds data security, integrity, and compliance while enhancing fine-grained governance capabilities within the data space.

This document defines the data transaction permission fields within the Data Fencing architecture and the whitelist information maintained by the gateway devices. By comparing the data transaction permission fields against the whitelist, data packets are either permitted or blocked, establishing precise data flow boundaries and egress control. Finally, the document presents several use cases of the data fencing.

2. Conventions and Definitions

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. Abbreviations and definitions used in this document:

*HBH: Hop-by-Hop Options Header.

3. Data Fencing

3.1. Data Fencing Architecture

In the data fencing architecture, the primary components include the Data Circulation and Utilization Platform, the Network Management Platform, Gateway Devices, and Access Devices. Access devices are connected to their respective gateway device, and different access devices can transmit data via links between gateway deivices.

Figure 1 shows the data fencing architecture.

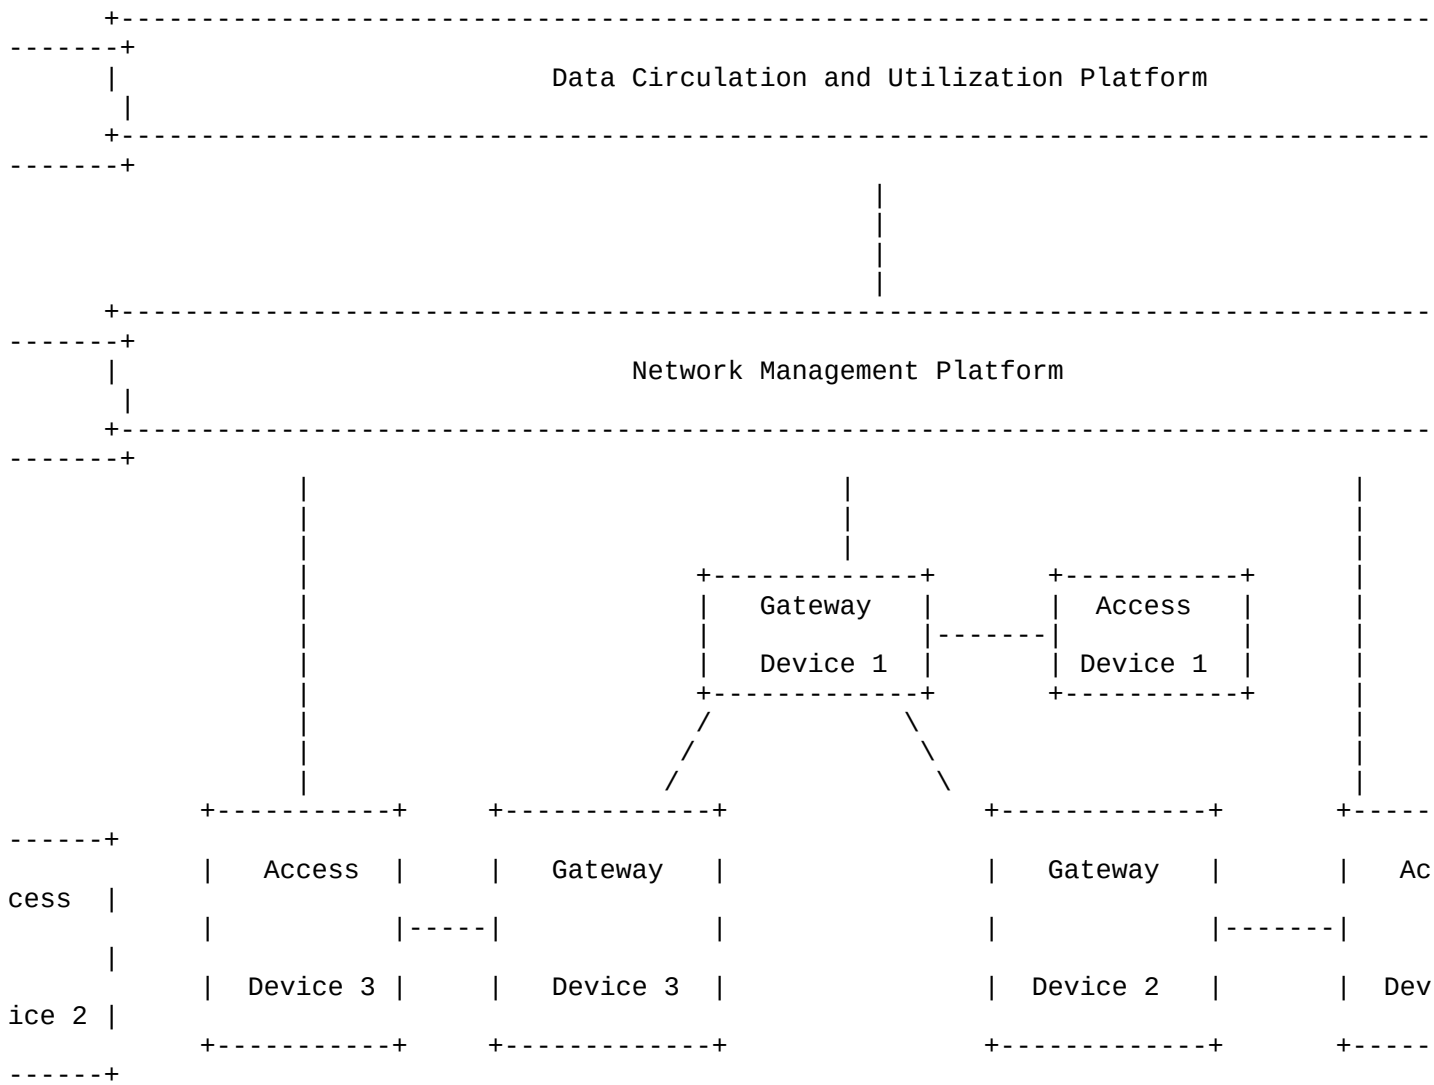

Figure 1: The data fencing architecture

The Data Circulation and Utilization Platform serves as a centralized marketplace that facilitates data transactions between data providers and consumers, aiming to break down data silos and promote seamless data exchange and sharing. In scenarios where certain users engage in frequent data transactions and seek to establish long-term, stable data sharing relationships, the platform offers specialized services through "Data Transaction Groups." Members within a Data Transaction Group are mutually trusted regarding data circulated inside the group, enabling secure and efficient collaboration. Each Data

Transaction Group is assigned a unique identifier for management and recognition.The Data Circulation and Utilization Platform sends Data Transaction Group information to the Network Management Platform.

The Network Management Platform is responsible for the flexible orchestration and scheduling of network resources. It supports the deployment of policies to network devices, enabling on-demand isolation and connection of access nodes as required. Additionally, this platform manages the unified allocation and administration of network identifiers, ensuring orderly and controlled network operations.The Network Management Platform deliver the Data Transaction Group information to the corresponding gateway devices.

The gateway device, serving as the network entry point, is capable of receiving IPv6 packets from data access endpoints. It maintains a dynamic whitelist in real-time, which constitutes a collection of authorized data group identifiers permitted for transmission.

For a specific data circulation instance, consider User 1 and User 2 as the authorized parties in the data transaction, while User 3 is not authorized. When data is being transmitted between User 1 and User 2, if the direct link between them experiences issues such as interruption or congestion, data packets may be rerouted along an alternative path, for example: User 1 -> User 3 -> User 2. Since User 3 is not authorized to access the data, it is imperative to prevent User 3 from gaining any visibility or access to the data packets during this process.

3.2. Data Transaction Permission Attribute TLV

The Data Transaction Permission Attribute can be directly encapsulated as an option in IPv6 extension headers within Hop-by-Hop Options Header(HBH) [RFC8200]. It’s a newly defined TLV used to to indicate whether the gateway device is permitted to forward the IPv6 packet.

Figure 2 shows the Data Permission Attribute TLV format.

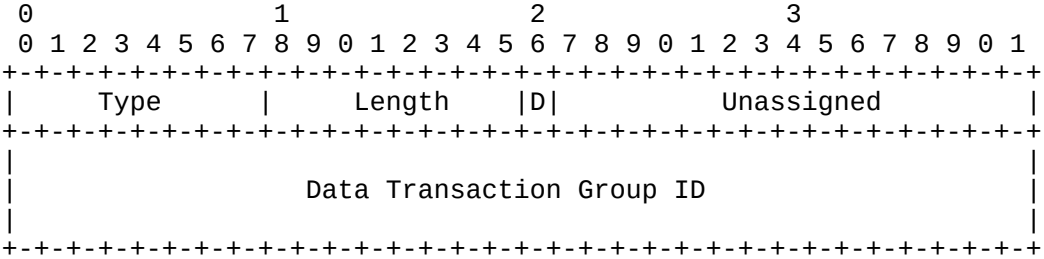

Figure 2: Data Permission Attribute TLV format

* Type: 8-bit identifier.

* Length: 8-bit unsigned integer that indicates the total number of the octets of the value field.

* D: 1-bit field used to indicate whether the current data transaction is directed.

* Unassigned: 15-bit field reserved for future use. They MUST be set to 0 on transmission and MUST be ignored on receipt.

* Data Transaction Group ID: 64-bit group ID of Data Transaction Group.

The Data Transaction Group ID is used to store the data transaction group to which the data belongs. Under normal circumstances, every member within a data trading group is considered trustworthy for data circulated inside the group.

The D field is used to indicate whether the current data transaction is directed. This field supports two values: 0 and 1.

When the field is set to 0, it means disabling directed data transaction, representing a standard data transaction where all members within the data transaction group are considered trustworthy and may access the group's data.

When the field is set to 1, it means enabling directed data transaction, representing a directed data transaction where the data is only accessible to specific designated users within the data transaction group. The information of these authorized users is delivered by the Network Management Platform to the relevant gateway devices for enforcement.

3.3. Whitelist

The gateway device maintains a whitelist of connected data access devices. This whitelist records information about devices or users permitted to access data through the gateway device. It contains two types of information: a Data Transaction Group ID and an Enable/ Disable Directed Data Transaction flag. These two types of information are updated after the gateway device receives instructions issued by the Network Management Platform. The main scenarios are as follows:

1. Data Transaction Group ID Update Updates to the Data Transaction Group ID generally occur when new users join the Data Circulation and Utilization Platform.

If a new user wishes to join one or more existing data transaction groups, the gateway device used by that new user will add the corresponding Data Transaction Group ID(s) to its whitelist.If a new user intends to establish one or more new data transaction groups, the Network Management Platform will notify all members within the group. Each member’s gateway device will then add the ID of the new data transaction group to its respective whitelist.

1. Enable/Disable Directed Data Transaction Update Updates to the Enable/Disable Directed Data Transaction flag(D field) generally occur during directed data transactions.By default, directed data transactions are disabled. When a directed data transaction is initiated, the initiator informs the Data Circulation and Utilization Platform of the intended participants. The Network Management Platform then notifies the gateway devices of the corresponding users, and these gateway devices set the D field to 1.Upon completion of the data transaction, the Network Management Platform again notifies the relevant gateway devices, which reset the D field to 0.

The gateway device can identify the information in each field of the network identifier. By comparing the Data Transaction Permission Attribute with the whitelist-stored information on permitted data groups, a determination is made regarding whether to discard the data packet, which means the gateway device conditionally allows or blocks data transmission. In the event of data being blocked, the gateway device reports a warning to the Network Management Platform, indicating a case of non-compliant data flow, thereby facilitating timely adjustment of data circulation paths and strategies.

4. Implementation Process of the Data Fencing Function

4.1. Information Registration of Data Transaction Parties

Upon initial access to the Data Circulation and Utilization Platform, users register by submitting their profile information, including specifications of the data trading groups they intend to join or establish.The Data Circulation and Utilization Platform updates the Data Transaction Group Information Table: If the user wishes to join one or more existing data transaction groups, the platform adds the user to the corresponding data trading group and updates the member information of the corresponding group(s). If the user wishes to establish one or more new data transaction groups, the platform creates the corresponding data trading group and updates the

information of the new group(s), including assigning a unique identifier to the new group and recording the member information within the new group. Then, the Data Circulation and Utilization Platform distributes the Data Transaction Group Information Table to the Network Management Platform via a synchronization interface.

Upon receiving messages from the Data Circulation Platform, the Network Management Platform distributes the Data Trading Group Information Table to gateway devices. If the user wishes to join an existing data trading group, the Network Management Platform delivers the identifier of the target data trading group to the user’s gateway device.If the user wishes to establish a new data trading group, the Network Management Platform delivers the new group information to the gateway device of all members (including the requesting user). The relevant gateway devices update their whitelists accordingly.

4.2. Data Transaction Process

Steps for Implementing Data Fencing via Whitelist on Gateway Devices for a Single Data Transaction

The access device on the provider side adds a Data Transaction Permission attribute to the data packet. This attribute carries information about the transaction, including the Data Transaction Group ID and whether it is a Directed Data Transaction. The packet is then sent to the provider’s gateway device. The gateway device on the data provider side forwards data packets to the corresponding network path based on policies delivered by the Network Management Platform.

Each gateway device along the transmission path inspects the Data Transaction Permission attribute and makes a forwarding decision as follows:The gateway device checks the Data Transaction Group ID field.If the group ID is not in the gateway’s whitelist, the packet is blocked (discarded).If the group ID is in the whitelist, the gateway device proceeds to inspect the D field.The gateway device then checks the D field:If the value is 0 (standard transaction), the packet is allowed.If the value is 1 (directed transaction), the gateway device checks whether directed transmission is enabled locally:If enabled -> the packet is allowed.If not enabled -> the packet is blocked (discarded).

If a gateway device blocks a packet, it reports a warning message to the Network Management Platform, including the device ID and relevant information about the discarded packet.

5. Use Cases

5.1. Cross-Domain Data Cooperation and Development

Data is increasingly becoming a core factor of production that drives innovation and growth. Individual organizations often face limitations in their internal data and capabilities, making it difficult to tackle complex development challenges or seize market opportunities alone. Therefore, it is necessary to integrate complementary data resources, specialized knowledge, and technical expertise from different institutions such as enterprises and research institutes. Through cross-domain data collaboration and development, organizations can break down information silos, leverage complementary strengths, and achieve mutual benefits through cooperation.

In cross-domain scenarios, it is essential to ensure that sensitive data shared between parties is used only for specific projects, by specific personnel, and within specific boundaries, preventing data from being copied, retained, or utilized for unauthorized purposes. Data fencing can establish a dynamic, mandatory security boundary that accompanies the data throughout its entire lifecycle. Its key functions include:

1. Establishing a Dynamic and Trusted Access Boundary: Only devices that are formally registered and authorized for participation in the project are included in the whitelist for data access. This ensures that data can only be accessed by predefined, trusted devices and users, effectively preventing unauthorized access at the point of entry.

2. Providing Real-Time Violation Alerting and Response Capabilities: Data fencing acts not only as a protective shield for data but also as an active defense system with monitoring and response capabilities. Leveraging data fencing functionality, when a gateway device detects non-compliant behavior such as unauthorized access or operation attempts, it can immediately report an alert to the management platform. The alert may include detailed information such as the identity of the violator, the time of the incident, and the specific nature of the violation. This shifts security protection from post-event analysis to real-time blocking and alerting during incidents, effectively guarding against both internal risks and external infiltration attempts.

5.2. Data Sharing and Isolation in Multi-Cloud Environments

Amid the wave of enterprise digital transformation, companies often adopt hybrid or multi-cloud architectures to optimize costs and leverage the unique advantages of different cloud service providers. As a result, business data naturally resides across diverse cloud platforms. For instance, core transaction data may be stored in a private cloud, while data requiring large-scale computational analysis is deployed in a public cloud. At the same time, cross-departmental collaboration and industrial chain coordination require data to flow securely within authorized boundaries. This demand for data sharing must be balanced against stringent data security and compliance requirements. Sensitive data, such as user personal information, financial records, and trade secrets, must be strictly isolated to prevent unauthorized access and data leaks, thereby meeting increasingly rigorous data protection regulations.

Data fencing addresses these challenges by establishing intelligent, policy-driven micro-boundaries around the data itself, ensuring secure and seamless data flow within authorized limits, even in complex multi-cloud environments.

1. Fine-Grained Data-Level Isolation and Policy Attachment: In hybrid cloud environments, data fencing moves beyond reliance on physical or network perimeters by directly binding security policies to data assets. This enables precise isolation of data with varying sensitivity levels even within the same database or the same table.

2. Unified Security Management to Reduce Multi-Cloud Complexity: Data fencing provides a unified security management plane, offering enterprises a holistic view and centralized control over data flows across clouds. This significantly simplifies the implementation of consistent security and compliance policies within heterogeneous IT environments.

5.3. Secure Data Acquisition and Boundary Control in IoT Environments

As production lines deploy a vast number of heterogeneous sensors and devices, massive volumes of equipment status, operational parameters, and production data are generated in real time. This data not only serves as a core asset for enabling predictive maintenance, process optimization, and intelligent decision-making but also constitutes business-critical information directly tied to production continuity and product quality. However, the inherent characteristics of IoT environments resource-constrained terminal devices, diverse communication protocols, and blurred network boundaries expose them to multiple security threats. These include unauthorized device

access, data tampering or theft during transmission, and the leakage of sensitive production data to unauthorized systems.

Data fencing establishes a micro-security fortress for IoT environments through identity authentication, integrity verification, and strict egress policies. This ensures that every piece of data flowing into the enterprise's digital core is trustworthy, intact, and controlled, laying a solid security foundation for truly intelligent industrial production.

1. Establishing a Trusted Identity Boundary and Access Control:Data fencing implements a strict whitelist mechanism at the collection gateway device (device gateway). The unique identifiers (e.g., MAC addresses, digital certificates) of every legitimate sensor and device are pre-registered in the data fencing whitelist during factory setup or deployment. Any unauthorized device attempting to connect is identified and blocked at the connection stage, thereby ensuring the purity and trustworthiness of data sources.

2. Constructing Precise Data Flow Boundaries and Egress Control:Data fencing defines the only legitimate exit points for data at the IoT gateway device, strictly controlling where data can be transmitted. Collected data can only be sent to one or several designated and authorized central analysis platform IP addresses. Any attempt to copy or forward data to unauthorized IP addresses is blocked in real time, preventing illegal data replication and leakage while ensuring that operational data remains within the enterprise's authorized scope of control.

6. Security Considerations

TBD

7. IANA Considerations

TBD

8. References

8.1. 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/info/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/info/rfc8174>.

8.2. Informative References

[RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <https://www.rfc-editor.org/info/rfc8200>.

Authors' Addresses

Xinrui Gu (editor)
China Unicom
Beijing
China
Email: guxr12@chinaunicom.cn

Naihan Zhang (editor)
China Unicom
Beijing
China
Email: zhangnh12@chinaunicom.cn

Xinxin Yi (editor)
China Unicom
Beijing
China
Email: yixx3@chinaunicom.cn