

FANTEL
Internet-Draft
Intended status: Informational
Expires: 23 April 2026

F. Zhang
J. Hu
Z. Hu
Y. Zhu
China Telecom
20 October 2025

FANTEL Use Cases and Requirements in Wide Area Network
draft-hhz-fantel-sar-wan-01

Abstract

This document introduces the main scenarios related to AI services in WAN, as well as their requirements for FANTEL (FAst Notification for Traffic Engineering and Load balancing) in these scenarios. Traditional network management mechanisms are often constrained by slow feedback and high overhead, limiting their ability to react quickly to sudden link failures, congestion, or load imbalances. Therefore, these AI services need FANTEL to provide real-time and proactive notifications for traffic engineering and load balancing, meeting the requirements of ultra-high throughput and lossless data transmission.

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 23 April 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. Conventions Used in This Document	3
2.1. Requirements Language	3
2.2. Abbreviations	3
3. Use Cases	4
3.1. AI Service Scenarios in WAN	4
3.1.1. Scenario 1: Sample Data Transfer	5
3.1.2. Scenario 2: Coordinated Model Training	5
3.1.3. Scenario 3: Coordinated Model Inference	6
3.2. Sample Data Migration	6
3.2.1. Use Case Description	6
3.2.2. Fast Notification Impact	6
3.2.3. Example	7
3.3. Remote Sample Data Access	8
3.3.1. Use Case Description	8
3.3.2. Fast Notification Impact	8
3.3.3. Example	8
3.4. Coordinated Model Training across AIDCs	9
3.4.1. Use Case Description	9
3.4.2. Fast Notification Impact	9
3.4.3. Example	9
3.5. Coordinated Model Training between Entities and AIDCs	10
3.5.1. Use Case Description	10
3.5.2. Fast Notification Impact	11
3.5.3. Example	11
3.6. Coordinated Model Inference	12
3.6.1. Use Case Description	12
3.6.2. Fast Notification Impact	12
3.6.3. Example	12
4. Challenges and Requirements	13
5. IANA Considerations	14
6. Security Considerations	14
7. References	14
7.1. Normative References	14
Contributors	15
Authors' Addresses	15

1. Introduction

The rapid development of Artificial Intelligence (AI), particularly large language models (LLMs), necessitates substantial computing power. Leasing computing resources from third-party AI Data Centers (AIDCs) provides a cost-efficient and elastic solution for entities such as industry enterprises or research institutions that find building and maintaining their own DC costly. However, AI service traffic, characterized by massive volume, high burstiness, and sensitivity to packet loss and latency, poses significant challenges to IP WAN interconnecting multiple entities and AIDCs. Moreover, entities with strict security requirements may prefer to keep their datasets on-premises, which introduces challenges in remote access and distributed coordination.

This document categorizes AI service scenarios over the WAN and analyzes representative use cases, including sample data migration, remote data access, coordinated model training between entities and AIDCs, coordinated model training across AIDCs, and coordinated model inference. Based on these use cases, this document summarizes the corresponding challenges and requirements, and discusses what and how the FANTEL architecture can address them.

2. Conventions Used in This Document

2.1. 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.

2.2. Abbreviations

- * AI: Artificial Intelligence
- * AIDC: AI Data Center
- * ECMP: Equal-Cost Multi Path
- * FANTEL: Fast Notification for Traffic Engineering and Load Balancing
- * INT: Inband Network Telemetry
- * IOAM: In-situ Operations, Administration, and Maintenance

- * LLM: Large Language Model
- * MAN: Metropolitan Area Networks
- * RDMA: Remote Direct Memory Access
- * TE: Traffic Engineering
- * TPOT: Time per Output Token
- * TTFT: Time to First Token
- * WAN: Wide Area Network

3. Use Cases

3.1. AI Service Scenarios in WAN

With the rapid growth of AI service traffic, entities and AIDCs face increasing strain on their computing resources. To address this, WAN provides the essential foundation for integrating and delivering computing power across sites. Based on WAN, on-demand compute leasing offers an elastic, cost-effective way to scale AI resources.

AI service scenarios in WAN can be categorized in into 3 scenarios, including sample data transfer, coordinated training, and coordinated inference, as shown in Figure 1.

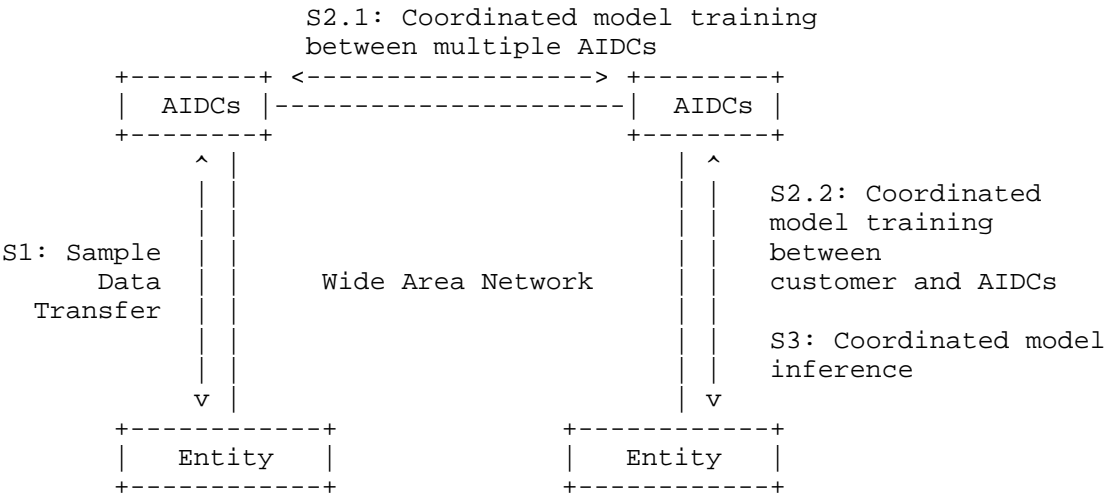


Figure 1: AI service scenarios in WAN

3.1.1. Scenario 1: Sample Data Transfer

Sample data transfer refers to the transfer of massive sample data from entities' storage DCs to their own or third-party AIDCs for model training. Due to the diversity of data security requirements vary among entities, there are two sub-scenarios:

- 1) Sample Data Migration: To meet the high-throughput and low-latency requirements of AI training, entities generally migrate their sample datasets to AIDC storages, where each training round can access terabytes to petabytes of data efficiently.
- 2) Remote Sample Data Access: To protect sensitive data, entities with strict security requirements retain their datasets in on-premises storage rather than migrating them to AIDCs. This creates the need for secure, low-latency, lossless approaches that allow timely remote access to sample data during model training.

3.1.2. Scenario 2: Coordinated Model Training

The computing power required for AI model training is enormous and grows rapidly, especially as models scale in size and complexity. For example, the computing power demand of LLMs grows rapidly — it is estimated that GPT-6 requires ZFLOPS-scale computing power, reaching a ~2000x increase over GPT-4. A single DC would struggle to meet such enormous demands. Therefore, integrating dispersed computing resources to support LLMs training (including training and fine-tuning of foundational models) has become a key solution.

- 1) Coordinated Model Training across AIDCs: The scalability of a single AIDC is inherently constrained by its physical infrastructure (e.g., space and power supply). In order to meet the enormous computing power requirements of training, one solution is to coordinate distributed computing resources across multiple AIDCs. It also helps fully utilize the idle computing resources available in different DCs.
- 2) Coordinated Model Training between Entities and AIDCs: Some entities that have deployed AI facilities in their own DCs have to face rapidly growing computing resource demands. Instead of build new AI infrastructure, which is costly, they can coordinate with leased third-party AIDCs to supplement their capacity. Moreover, entities with strict security requirements can further enhance this approach by incorporating split learning with the input/output layers deployed locally to ensure sensitive data remains on-premises.

3.1.3. Scenario 3: Coordinated Model Inference

Entities that have deployed AI facilities in their own DCs also have to face the degradation of TTFT and TPOT as inference concurrency increases continuously. Since expanding on-premises AI capacity is often costly, leasing third-party AIDCs and coordinating inference between entities' AIDCs and third-party AIDCs can be a cost-effective way to scale concurrent inference capacity.

3.2. Sample Data Migration

3.2.1. Use Case Description

Since AI training requires multiple rounds of fine-tuning for performance improvement, and each round consumes massive sample data (ranging from terabytes to petabytes), entities need to upload these sample data into AIDCs as soon as possible to start the next round of training.

Currently, many entities still rely on shipping physical hard drives to migrate such large datasets, which not only risks data loss if the drives are damaged but is also highly inefficient.

For network-based solutions, entities typically have to rent dedicated line services with fixed-bandwidth on a monthly or annually subscription basis, which is less cost-effective because the data transfer traffic is bursty, meaning the bandwidth is only fully utilized during short transfer periods and remains idle for the rest of the time.

3.2.2. Fast Notification Impact

For high-efficiency and cost-effective sample-data migration, fast notification of network status enhances the WAN in the following ways:

- 1) To maximize the available bandwidth for the data transmission, WAN should support fast notification of network status changes across devices, achieving efficient hourly transmission of terabyte-scale sample data by enabling real-time service bandwidth adjustment based on tasks. Moreover, the data migration services should be provisioned on a task basis, eliminating the need for entities to lease high-bandwidth lines on a monthly/yearly basis and thereby significantly reducing costs.

2) During high-speed sample data migration, even minor network failures can trigger significant packet loss, sharply reducing the efficiency. To prevent this the WAN should provide millisecond-level fast failure notification, enabling rapid failure detection and failover to maintain high throughput.

3.2.3. Example

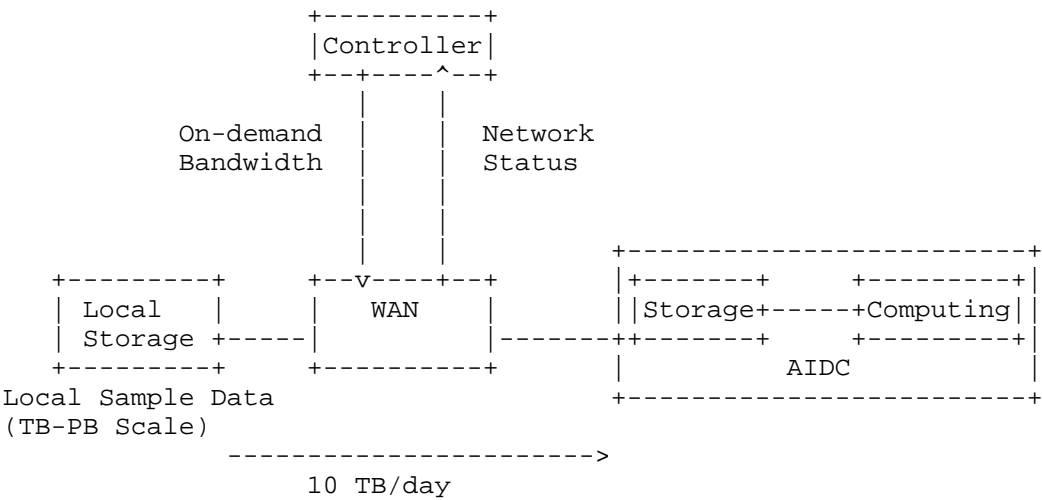


Figure 2: Sample Data Migration from Local Storage to AIDC

An example of sample data migration is an enterprise that leases the AI services from a third-party AIDC. In this case, the enterprise collects and stores the sample data in its local storage and needs to transfer 10 TB data to AIDC every day. It would take several days if it transferred by shipping hard drives, or 10 days if transfer over a 100 Mbps link.

Using real-time network status obtained from the fast notification mechanism, the controller can flexibly adjust service bandwidth on-demand in seconds and ensure high throughput via traffic engineering and load balancing. In this example, the controller adjusts the service bandwidth to 10 Gbps for 3 hours, which is sufficient to complete the 10 TB data migration, and dynamically updates the traffic engineering and load balancing policies to maintain high throughput.

Zhang, et al. Expires 23 April 2026 [Page 7]

3.3. Remote Sample Data Access

3.3.1. Use Case Description

Some industries with highly sensitive data prefer to keep their datasets on-premises, avoiding the risk of leakage that may arise when migrating data to third-party AIDCs.

The most common method for accessing such data during AI computing is to use RDMA protocols (e.g. InfiniBand and RoCE), which achieve ultra-low latency, actively bypass TCP's congestion control mechanism, and rely on the Go-back-N mechanism to handle packet loss and disorder. The Go-back-N mechanism retransmits all unacknowledged packets (including correctly received ones) after a timeout, causing the RDMA protocol extremely sensitive to packet loss -- even a 0.1% packet loss can cause throughput to drop by roughly 50%.

To provide efficient AI services for these industries, robust congestion-control solutions are needed in WANs to minimize latency and packet loss.

3.3.2. Fast Notification Impact

Millisecond-level, lightweight notification can be sent to nodes adjacent to or affected by failure or congestion, enabling lossless transmission of sample data and meeting the stringent packet-loss tolerance requirement of RDMA transmission.

3.3.3. Example




Figure 3: Remote Sample Data Transmission via RDMA

An example of remote data access involves an enterprise with strict security requirements that prohibit storing sensitive data outside its premises, while still wishing to lease computing resources from

third-party AIDCs. In this case, remote sample data are transmitted between the AIDC and the enterprise's local storage via RDMA, which is highly sensitive to packet loss. The distance between the AIDC and the local storage may range from 100 to 500 km.

The fast notification mechanism enables flow-based precise congestion control through immediate congestion notification, ensuring lossless RDMA transmission and thereby supporting secure and efficient model training.

3.4. Coordinated Model Training across AIDCs

3.4.1. Use Case Description

Due to the limited computing resources of a single DC, the training task can be split and coordinated across multiple AIDCs. This approach also helps fully utilize the idle computing resources available in different DCs. Coordinated model training across AIDCs requires the WAN to support massive, highly concurrent and bursty traffic of parameter synchronizations.

3.4.2. Fast Notification Impact

For coordinated model training across AIDCs, fast failure and congestion notification enhances WAN in the following ways:

- 1) Dynamic load balancing: Fast network status notifications allow dynamic load balancing strategies to be deployed in real time, ensuring optimal utilization of network resources and maintaining high performance.
- 2) Low-latency, lossless parameter synchronization: Fast notification enables millisecond-level congestion control in the WAN, allowing upstream devices to promptly reduce their transmission rates upon detecting impending congestion.
- 3) Rapid failure protection: Interruptions in parameter synchronization due to network failures can trigger rollback and computation waste, sharply reducing training efficiency [draft-cheng-rtgwg-ai-network-reliability-problem]. Fast notification enables millisecond-level failure detection and failover, minimizing training disruptions.

3.4.3. Example




Figure 4: Coordinated Model Training across AIDCs

An example of coordinated Model Training across AIDCs is splitting the training task of LLM using parallelization strategies such as pipeline parallelism and data parallelism. During model training, LLM parameters are synchronized across geographically distributed AIDCs via the WAN using the RDMA protocol. The parameter synchronization traffic is highly concurrent and bursty elephant flows, characterized by long duration and large data amount, which can easily cause network congestion.

Leveraging fast network status notification, the WAN can perform efficient traffic engineering and load balancing. In addition, fast failure and congestion notifications enable flow-based precise congestion control, ensuring lossless and efficient parameter synchronization.

3.5. Coordinated Model Training between Entities and AIDCs

3.5.1. Use Case Description

Considering cost and security, some entities may choose to lease third-party AIDCs to meet their rapidly growing computing resource demands. In this case, split learning can be applied, where only input/output layers are deployed locally for data security, while the intermediate layers are deployed in third-party AIDCs for cost efficiency. Activations and gradients are transmitted via the WAN. However, the transmission between entities and AIDCs still requires low latency and even more elastic bandwidth.

3.5.2. Fast Notification Impact

For coordinated model training between entities and AIDCs, fast notification mechanism enhances WAN performance in the following ways:

- 1) Providing low latency through fast congestion and failure notification (as discussed in Section 3.4.2).
- 2) Enabling elastic bandwidth allocation and dynamic traffic engineering and load balancing strategies through fast network status notification (as discussed in Section 3.2.2).

3.5.3. Example

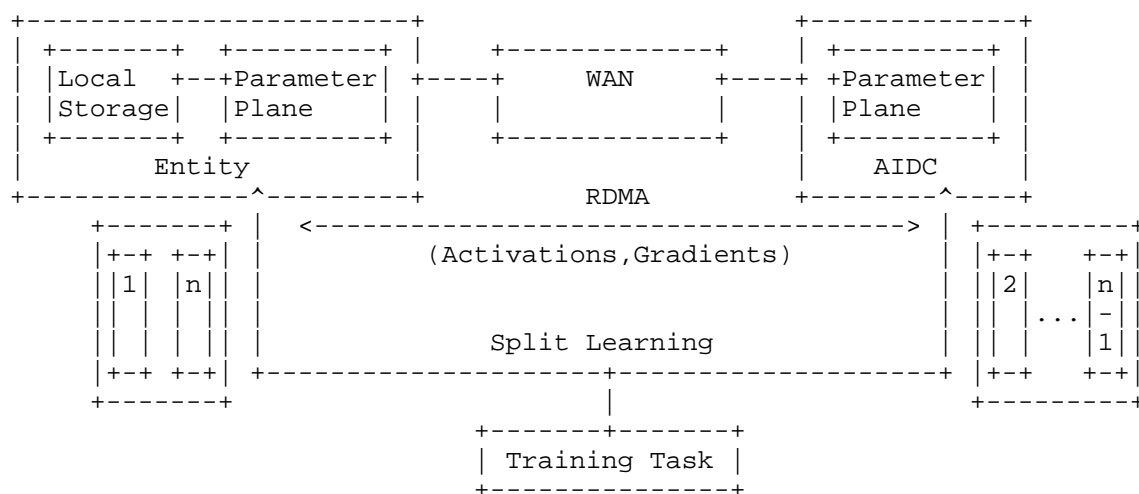


Figure 5: Coordinated Model Training between entities and AIDCs

An example of coordinated model training between entities and AIDCs is that an entity only has to build minimum amount of computing resource to deploy input/output layers locally, while leasing third-party AIDC resources to deploy intermediate layers. During model training, the activations of the forward pass and gradients of the backward pass are transmitted via the WAN using the RDMA protocol, which are bursty and latency-sensitive.

Fast network status notification enables flexible bandwidth allocation, dynamic traffic engineering and load balancing, while fast congestion notification ensures low latency.

3.6. Coordinated Model Inference

3.6.1. Use Case Description

Similar to Section 3.5, entities may also choose to lease third-party AIDCs for model inference. Some may completely depend on third-party AIDC for inference which require enough bandwidth and low latency to access via the WAN, while others lease third-party AIDCs as the supplement coordinating with the entities' local inference which requires low latency, low packet loss and cost-effective transmission.

3.6.2. Fast Notification Impact

For coordinated model inference, fast notification mechanism enhances WAN performance in the following ways:

- 1) Providing low latency through fast congestion and failure notification (as discussed in Section 3.4.2).
- 2) Enabling elastic bandwidth allocation and dynamic traffic engineering and load balancing strategies through fast network status notification (as discussed in Section 3.2.2).

3.6.3. Example

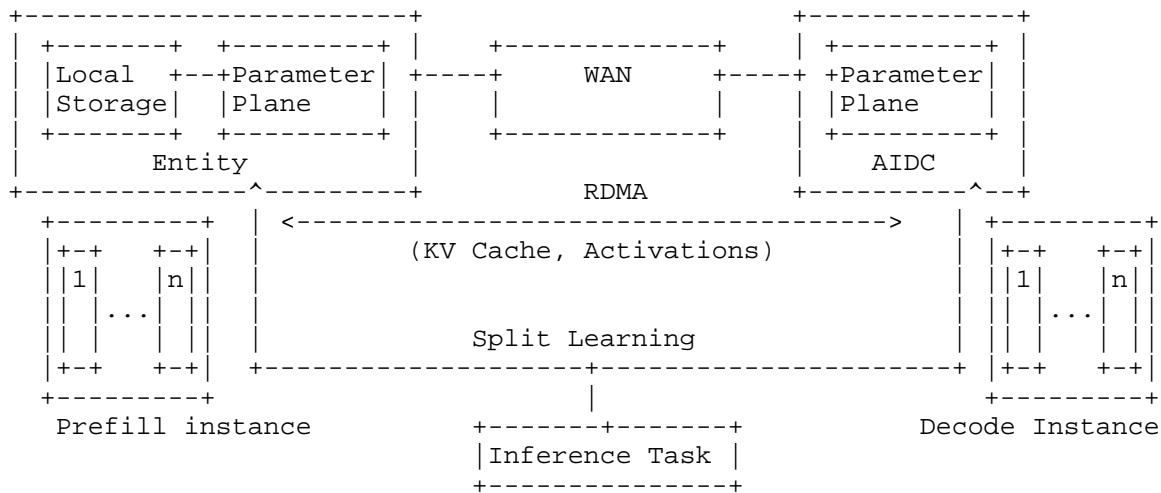


Figure 6: Coordinated Model Inference

An example of coordinated model inference employs split learning to distribute the inference task between entities and AIDCs. The prefill instance is deployed locally and the decode instance is deployed in third-party AIDCs, since the decode phase has significantly higher GPU memory requirements compared to the prefill phase. This reduces the demand on the entity's local DC and keeps the prompts on-premises. Additionally, the input and output layers of the decode phase can remain in the entity's DC to meet stricter data security requirements. During model inference, the key/value cache and intermediate activations are transmitted via the WAN using the RDMA protocol.

Similar to Section 3.5, fast network status notification enables flexible bandwidth allocation, dynamic traffic engineering and load balancing, while fast congestion and failure notification ensures low latency and effective congestion control.

4. Challenges and Requirements

The above use cases introduce elastic, cost-effective, and secure ways to scale and fully utilize the computing resources for AI services across multiple sites. However, these approaches involve long-distance transmission over the WAN of AI service traffic, which is characterized by massive volume, high burstiness, and sensitivity to packet loss and latency. These characteristics expose limitations in existing mechanisms, including delayed decision-making, coarse-grained feedback, and slow recovery:

- * Load Balancing mechanisms (e.g., IOAM) typically depend on centralized control or static policies, resulting in delayed reaction to highly dynamic traffic, which may lead to congestion or packet-loss.
- * Flow Control mechanisms (e.g., ECN [RFC3168]) often rely on end-to-end feedback that is constrained by RTT delays, making it hard to achieve fine-grained, real-time adjustment.
- * Failure Protection mechanisms (e.g., BFD [RFC5880], FRR [RFC7490]) typically involve periodic detection and precomputed backup paths, which cannot always provide millisecond-level recovery in complex multi-domain environments. Moreover, increasing probe frequency to shorten detection time inevitably raises CPU and bandwidth overhead.

To address these challenges, FANTEL architecture needs to provide fast, real-time, lightweight notifications for efficient load balancing, flow control, and failure protection, enabling elastic bandwidth, lossless transmission, and fast failover recovery for AI service traffic over the WAN, including:

- * Fast Network Status Notification delivers real-time visibility into traffic patterns, link utilization, and node load to support timely adjustments of paths and traffic rates.
- * Fast Congestion Notification provides low-latency, fine-grained feedback to enable immediate adjustments of data transmission rate or re-routing, preventing congestion and packet loss.
- * Fast Failure Notification notifies link or node failures with real-time detection and precise propagation, allowing immediate responses such as switching to backup paths, rerouting traffic, or suppressing affected routes to ensure service reliability.

5. IANA Considerations

N/A.

6. Security Considerations

TBD.

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/info/rfc2119>>.
- [RFC3168] Ramakrishnan, K., Floyd, S., and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP", RFC 3168, DOI 10.17487/RFC3168, September 2001, <<https://www.rfc-editor.org/info/rfc3168>>.
- [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>>.

- [RFC7490] Bryant, S., Filsfils, C., Previdi, S., Shand, M., and N. So, "Remote Loop-Free Alternate (LFA) Fast Reroute (FRR)", RFC 7490, DOI 10.17487/RFC7490, April 2015, <<https://www.rfc-editor.org/info/rfc7490>>.
- [RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", RFC 5880, DOI 10.17487/RFC5880, June 2010, <<https://www.rfc-editor.org/info/rfc5880>>.
- [draft-cheng-rtgwg-ai-network-reliability-problem]
"Gap Analysis of Fast Notification for Traffic Engineering and Load Balancing".

Contributors

Thanks to all the contributors.

Authors' Addresses

Fan Zhang
China Telecom
109, West Zhongshan Road, Tianhe District
Guangzhou
Guangdong, 510000
China
Email: zhangf52@chinatelecom.cn

Jiayuan Hu
China Telecom
109, West Zhongshan Road, Tianhe District
Guangzhou
Guangdong, 510000
China
Email: hujy5@chinatelecom.cn

Zehua Hu
China Telecom
109, West Zhongshan Road, Tianhe District
Guangzhou
Guangdong, 510000
China
Email: huzh2@chinatelecom.cn

Yongqing Zhu
China Telecom
109, West Zhongshan Road, Tianhe District
Guangzhou
Guangdong, 510000
China
Email: zhuyq8@chinatelecom.cn