

FANTEL
Internet-Draft
Intended status: Informational
Expires: 7 January 2026

J. Hu
Z. Hu
Y. Zhu
China Telecom
6 July 2025

FANTEL scenarios and requirements in Wide Area Network
draft-hhz-fantel-sar-wan-00

Abstract

This document introduces the main scenarios related to AI services in WAN, as well as the 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 new AI services need FANTEL to provide real-time and proactive notifications for traffic engineering and load balancing, meeting the ultra-high throughput and lossless data transmission requirements of these AI service scenarios.

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 7 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	2
2. Conventions Used in This Document	3
2.1. Requirements Language	3
2.2. Abbreviations	3
3. Use Cases	4
3.1. Scenario 1: Sample data transmission	4
3.1.1. Sub-scenario 1.1: Transmitting sample data into storage system	4
3.1.2. Sub-scenario 1.2: Directly transmitting sample data to AI servers	5
3.2. Scenario 2: Coordinated model training	5
3.3. Scenario 3: Coordinated model inference	6
4. Problem Statement	7
5. Requirements	8
6. IANA Considerations	9
7. Security Considerations	9
8. References	9
8.1. Normative References	9
Contributors	10
Authors' Addresses	10

1. Introduction

The rapid development of Large Language Models (LLMs) necessitates substantial computing power. Hyperscalers build their own AI Data Centers (AIDCs) to train foundational models. However, most enterprises do not have a demand for training foundational models but want to meet the needs of fine-tuning and inference cost-effectively, so a good solution is to rent third-party AIDCs for LLMs fine-tuning and inference, requiring IP network to fulfill their needs. IP network consists of IP Backbone and IP Metropolitan Area Networks (IP MAN). IP MAN interconnects various customers and data centers (including AIDCs) within the metropolitan area, while IP Backbone interconnects IP MANs and data centers (including AIDCs). IP Backbone and IP MAN belong to IP Wide Area Network (IP WAN, or WAN for short).

The AI services in WAN, including sample data transmission, coordinated model training and inference, require networks to efficiently manage traffic and rapidly adapt to network changes.

[draft-geng-fantel-fantel-requirements] points out that existing network management mechanisms such as FRR, BFD and ECN which often rely on delayed feedback or reactive responses, resulting in network performance degradation, longer service disruptions, or inefficient resource utilization. Therefore, FANTEL is proposed to implement real-time and reliable notifications of network events, effectively supporting Traffic Engineering (TE) functions such as load balancing, failure protection, and congestion control. WAN need to deploy FANTEL to ensure high throughput and lossless transmission of data, meeting the new demands of AI services.

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

AIDC: AI Data Center

ECMP: Equal Cost Multi Path

FANTEL: Fast Notification for Traffic Engineering and Load Balancing

INT: Inband Network Telemetry

LLM: Large Language Model

MAN: Metropolitan Area Networks

RDMA: Remote Direct Memory Access

TE: Traffic Engineering

WAN: Wide Area Network

3. Use Cases

For most customers, the cost of owning and maintaining AI facilities is prohibitively high. A solution is to rent AI facilities, which are located in third-party AIDCs to fulfill customers' LLMs training or inference requirements conveniently. Under these circumstances, customers access AIDCs via WAN to get various AI services as shown in Figure 1, including sample data transmission, coordinated model training across AIDCs, and coordinated model inference between the customer and AIDCs.

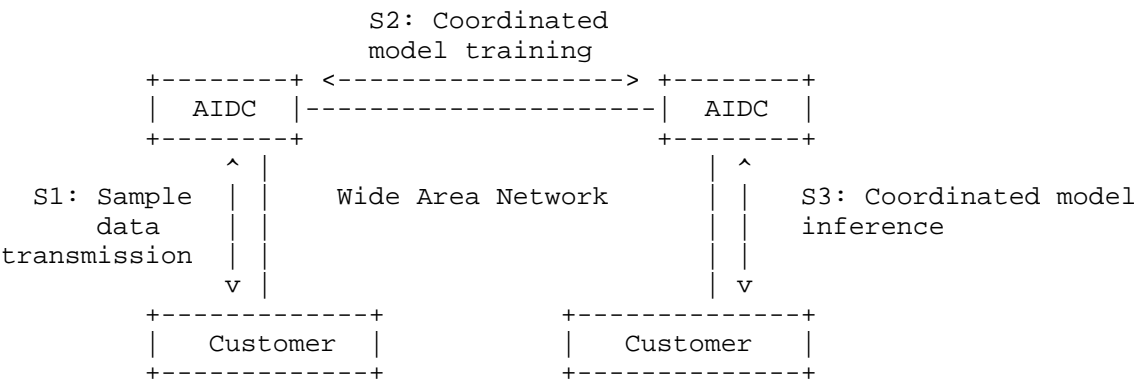


Figure 1: AI service scenarios in WAN

3.1. Scenario 1: Sample data transmission

When customers train AI models in third-party AIDCs, they need to transmit massive sample data into AIDCs. Due to differentiated customers' requirements for data security, there are two sub-scenarios involved.

3.1.1. Sub-scenario 1.1: Transmitting sample data into storage system

Since the training process of LLMs needs rounds of fine-tuning for performance improvement, with each round consuming massive sample data (ranging from terabytes to petabytes). These customers need to upload sample data into AIDCs as soon as possible to start the training. To provide high-efficiency and cost-effective sample data transmission, WAN needs to meet the following requirements:

1. Customers usually obtain fixed bandwidth based on dedicated line services to meet their service needs. However, the customer's sample data transmission requirements are intermittent and need to be as fast as possible, thus requiring WAN to have the ability to flexibly adjust bandwidth based on tasks.

2. To maximize the available bandwidth for the data transmission, this scenario requires the WAN to support fast notification of network status changes between devices, enabling real-time service bandwidth adjustment, enabling efficient hourly transmission of terabyte-scale sample data.

3. During high-speed sample data transmission, network failures can lead to a large number of packet losses, resulting in a sharp decrease of data transmission efficiency. This scenario requires WAN to have millisecond-level fast failure protection capability, enabling rapid failure detection and failover.

3.1.2. Sub-scenario 1.2: Directly transmitting sample data to AI servers

Certain customers with stringent data security requirements prohibit storing sample data outside their facilities. To address this problem, sample data must be uploaded to the AI training servers by using RDMA protocols while these servers are performing training tasks. Current mainstream RDMA protocols rely on Go-Back-N mechanism, making them highly sensitive to latency and packet loss (Even a 0.1% packet loss rate can degrade computational efficiency by 50%).

To achieve lossless transmission of sample data, based on the requirements of sub-scenario 1.1, WAN needs further support millisecond-level fast congestion control, meeting the requirement of zero packet loss in RDMA transmission. Therefore, this scenario requires the mechanism of FANTEL to quickly notify upstream devices to reduce traffic speed based on the router cache situation.

3.2. Scenario 2: Coordinated model training

The scaling laws demonstrate the performance of LLMs scales with model size, sample dataset size, and the amount of computing power used for training. The computing power demand of LLMs grows rapidly (It's estimated that GPT-6 requires ZFLOPS-scale computing power, reaching a ~2000x increase over GPT-4). The model training requirements of customers consist of foundational model training and fine-tuning. The computing power of a single AIDC is limited by physical infrastructure (e.g., space and power supply), making it inadequate to meet the demands of LLM training. Thus, a solution is proposed to fulfill the computing power demand of ultra-scale LLMs training through the efficient coordination of distributed computing resources across multiple AIDCs. Besides that, there are always some residual computing resources that are insufficient to meet the demands of a single customer. These resources can be coordinated across AIDCs to meet more customers' demands. For customers with

high data security requirements and self-built DCs, the collaborative training between the customer's own DC and AIDC can achieve ultra-scale LLMs training while ensuring sample data does not leave customer's DC (Requiring the input and output layers of LLM to be deployed within the customer's DC).

In this scenario, the training task of LLM is split across multiple AIDCs based on parallelization strategies such as pipeline parallelism and data parallelism. During model training, the parameters of LLMs need to be synchronized among AIDCs. The synchronization traffic of the parameter plane is transmitted via RDMA protocol which usually features some elephant flows. Therefore, WAN should provide efficient and lossless transmission of parameter plane data. WAN needs to meet the following requirements:

1. The characteristics of elephant flow are long duration and large data amount, which can easily cause network congestion. The synchronization of parameter plane requires low latency and zero packet loss. This scenario requires WAN to have millisecond-level fast congestion control capability, rapidly notifying upstream devices to slow down traffic rates upon detecting impending congestion. In addition, because traditional load balancing often relies on static policy, WAN needs to have a fast response for load balancing, immediately adjust load balancing decisions in response to network changes, ensuring optimal resource utilization and performance.

2. Interruption of parameter plane synchronization due to network failure may result in breakpoint rollback, causing wastage of computing power, leading to a sharp decrease in computational efficiency [draft-cheng-rtgwg-ai-network-reliability-problem]. This scenario requires WAN to implement millisecond-level failure protection, which can quickly detect network failures and failover.

3.3. Scenario 3: Coordinated model inference

Many customers have deployed AI servers in their own DCs to support LLM inference applications. However, the high deployment cost and operational complexity of on-premises deployment limit the scale of computing power. Due to the increasing inference concurrency, this on-premises deployment method cannot meet the computing power demand. To address this, the collaboration model inference between customer and AIDCs presents a more efficient, agile, and cost-effective approach to realize elastic computing power scaling.

In this scenario, the training task of LLM is split across customers and AIDCs based on parallelization strategies such as pipeline parallelism and expert parallelism. Taking the LLMs inference based

on Prefill-Decode disaggregation architecture as an example, the input and output layers of Prefill/decode are placed in the customer, while other layers are placed in the AIDC, ensuring large-scale inference concurrency, utilizing the computing resources in AIDCs to handle larger-scale inference concurrency and ensuring that sample data does not leave customer's DC.

During model inference, the parameter synchronization between AIDCs are transmitted via RDMA protocol. Similar to scenario 2, this scenario also requires WAN to have real-time elephant flow load balancing, millisecond-level congestion control, and fast network failure protection capabilities.

4. Problem Statement

According to the AI scenarios mentioned above, the primary challenge for WAN is real-time traffic engineering and load balancing. Current traffic engineering mechanisms have difficulty providing low-latency and low-overhead solutions that meet the above requirements, presenting the following issues:

1. Current load balancing techniques face great challenges in highly dynamic environments. One of the core issues is the lack of timely awareness and adaptive response to network state changes. Traditional mechanisms often rely on periodic global state synchronization or static policies, which results in delayed decision-making. The current controller-based load balancing uses In-situ OAM (IOAM) to obtain network status information. IOAM provides visibility into traffic by embedding telemetry data directly in packets. However, IOAM data is extracted and reported by the device CPU to a controller which adds latency and limits responsiveness [draft-geng-fantel-fantel-gap-analysis]. Moreover, controllers typically process telemetry in software, resulting in delayed decision-making. The delay of controller-based load balancing typically at second-scale, inevitably leads to network congestion and severe packet loss.

2. Existing flow control mechanisms rely on delayed feedback or reactive responses, which can lead to suboptimal network performance in high-latency or long-RTT environments like WAN. TCP-based congestion control is a receiver-driven congestion control which uses feedback signals from the receiver to adjust the transmission rate of the sender. These signals are subject to RTT delays, especially problematic in high-speed dynamic environments. ECN [RFC3168] marks packets to indicate congestion. However, ECN relies on end-to-end signaling and lacks precise real-time feedback. INT provides path-level telemetry by inserting metadata at each hop, which is returned to the sender via the ACK. Some congestion control algorithms, such

as High Precision Congestion Control (HPCC), utilize INT for precise load-awareness. The telemetry based on INT generates an RTT delay before the sender receives feedback, which limits the response capability. These end-to-end signaling-based flow control mechanisms introduce tens of milliseconds of latency in large-scale WAN, which fails to meet the requirements for lossless data transmission.

3. Existing failure protection mechanisms like BFD [RFC5880] and FRR [RFC7490] are widely deployed, they both have limitations in speed and scope. BFD is designed for rapid failure detection by sending frequent control packets between peers, but high probe frequency not only increases CPU and bandwidth usage but also strains the control plane in large-scale networks. Furthermore, 50ms detection cycle also makes it difficult for BFD to meet the detection requirements of some large-scale networks for link failures. Routing convergence mechanisms depend on routing protocol convergence, which may take hundreds of milliseconds. FRR serves as the complementary mechanism to routing convergence, achieving millisecond-level failover through pre-computed backup paths. Due to protecting against only adjacent failures, FRR lacks flexibility and responsiveness in complex topologies, with recovery latency reaching tens of milliseconds. Traditional Failure protection mechanisms rely on periodic failure detection and centralized rerouting, resulting in recovery times that are not fast enough.

5. Requirements

To solve the above-mentioned problems, FANTEL is needed to provide real-time, rapid notification of network events to relevant network nodes, including:

1. Fast network status notification. FANTEL uses traffic state detection to monitor traffic patterns, link utilization, and node load to trigger notifications on significant deviations [draft-geng-fantel-fantel-requirements]. Nodes can adjust the path and traffic rate in real-time based on FANTEL, achieving link status to achieve efficient traffic engineering and load balancing.

2. Fast congestion notification. FANTEL provides a fast, low-latency notification mechanism that can detect and alert network devices to congestion events in real time. When congestion occur, node can adjust data transmission rate and re-route the transmission route based on FANTEL, preventing packet loss.

3. Fast failure protection. FANTEL uses fast failure detection and notification to monitor real-time link/node status. When failure occurs, a node with protection mechanisms may immediately switch to backup paths, reroute traffic, or suppress affected routes, ensuring service reliability.

In summary, FANTEL provides a real-time notification mechanism that can be used in WAN, enabling bandwidth utilization, lossless transmission, and fast failover in different AI scenarios.

6. IANA Considerations

TBC

7. Security Considerations

TBC

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>>.
- [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-geng-fantel-fantel-requirements] "Requirements of Fast Notification for Traffic Engineering and Load Balancing".

[draft-geng-fantel-fantel-gap-analysis]

"Gap Analysis of Fast Notification for Traffic Engineering
and Load Balancing".

[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

Jiayuan Hu
China Telecom
109, West Zhongshan Road, Tianhe District
Guangzhou
Guangdong, 510000
China
Email: hujiy5@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