

Network Working Group
Internet-Draft
Intended status: Informational
Expires: 4 September 2025

J. Dong
Huawei Technologies
D. Li
Tsinghua University
Q. Shi
Huawei Technologies
P. Huo
ByteDance
3 March 2025

Current State of the Art for Routing in AI Networks
draft-dong-fantel-state-of-art-01

Abstract

This document provides an overview of routing technologies that address the needs of traffic engineering and load balancing, with a focus on fast notification for example in adaptive or perceptive routing. As the scale and complexity of networks grow, these technologies are becoming increasingly important when fault tolerance and rapid convergence are critical. This document explores existing solutions from both the IETF and the broader industry, highlighting their applicability to various use cases, including AI workloads and general services that demand low-latency, fault recovery, and dynamic load distribution across data center networks and inter data center. It also offers suggestions for potential IETF initiatives to further develop and standardize these techniques.

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 4 September 2025.

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. Proposals in IETF	3
2.1. Gap Analysis, Problem Statement and Requirements	3
2.2. Framework	4
2.3. Information Model	5
2.4. Solutions	5
2.4.1. Topology-specific Routing Mechanisms	5
2.4.2. Extensions to Routing Protocols	6
2.4.3. New Protocols for Fast Notification	6
3. Implementations in Industry	7
3.1. DLB and GLB	7
3.2. VRF-based Adaptive Routing	7
3.3. CONGA	8
3.4. Centralized TE and E-ECMP	8
4. Summary and Potential Work	8
5. Security Considerations	9
6. IANA Considerations	9
7. Acknowledgments	9
8. Informative References	9
Authors' Addresses	12

1. Introduction

The emergence of new applications, like AI applications brings new requirements to networks, such as load balancing and network reliability. AI-driven applications tend to generate highly dynamic and unpredictable traffic patterns, and require high performance in terms of throughput, latency and packet loss. As a result, there is a growing need for Adaptive and Perceptive Routing mechanisms that can respond to these new demands. As widely discussed both in standards and the industry, Adaptive/Perceptive Routing allows networks to make real-time adjustments in response to varying traffic

loads and network conditions, ensuring that network resources are optimally utilized and congestion is minimized.

This document provides an overview of routing technologies that address the needs of traffic engineering and load balancing, with a focus on the fast notification, for example in adaptive routing where the routing decision adapts to network events. As the scale and complexity of networks grow, these technologies are becoming increasingly important when fault tolerance and rapid convergence are critical. This document explores existing solutions from both the IETF and the broader industry, highlighting their applicability to various use cases, including AI workloads and general services that demand low-latency fault recovery and dynamic load distribution across data center networks and inter data center. It also offers suggestions for potential IETF initiatives to further develop and standardize these techniques.

2. Proposals in IETF

There are several individual drafts in the IETF which describe the problems, gaps, requirements and potential frameworks for routing in AI networks. This section briefly goes through these documents, summarizes the current state of this topic in the IETF, and identifies the open issues which needs further work.

2.1. Gap Analysis, Problem Statement and Requirements

[I-D.hcl-rtgwg-ai-network-problem] analyzes the gaps in the networks used for AI training, and describes the requirements for improvements. It firstly introduces the characteristics of AI training traffic, then focuses on the gaps and requirements in several key technologies: Load Balancing, Congestion Control and Fast Failover. It is not clear whether the congestion control mentioned in this document is more related to the network layer or the transport layer.

[I-D.cheng-rtgwg-ai-network-reliability-problem] focuses on the reliability problem and requirements in AI networks. It describes the existing mechanisms for network reliability, including link fault detection, ECMP, fast reroute and fast route convergence, (e.g. BGP Prefix Independent Convergence (PIC)), then analyzes the gaps in the timing of fault detection, notification propagation and switchover. In the end, the draft lists a set of requirements for new techniques on fault detection, congestion elimination, fast fault notification and fast switching over.

[I-D.wang-rtgwg-dragonfly-routing-problem] introduces the characteristics and routing mechanisms of dragonfly topology, including Minimal Routing, Non-Minimal Routing, Adaptive Routing and Valiant Load-Balanced Routing. Then it analyzes the gaps of existing routing mechanism in dragonfly networks, such as load balancing and adaptive routing notification, in the end the drafts list the requirements on routing protocols for dragonfly networks.

The analysis shows that there are some overlaps in the gap analysis and problem statement between these documents. The common problems and gaps identified for routing in AI networks are load balancing and fast failure notification. The requirements on routing protocols and the notification mechanism need further investigation.

2.2. Framework

[I-D.cheng-rtgwg-adaptive-routing-framework] describes a framework for adaptive routing, including a set of components, their interaction and the workflow. It identifies the problems with existing flow-based load balancing in AI networks, especially when congestion happens on some of the links. The solutions are classified into two types: flow-based adjustments and packet-based adjustments. The flow-based adjustments are further categorized into weight-based dynamic ECMP and Flow redirection. The overall adaptive routing framework consists of routing plane, forwarding plane, adaptive routing policy and the remote congestion detection. In the forwarding plane, it proposes to add remote path info to the forwarding table, and the quality of the links can be updated in response to congestion, then new weight value can be calculated to optimize the weight-based load balancing. In the routing plane, the draft analyzes the possible extensions needed in routing protocols for obtaining the path information. In congestion detection, it gives the definition of congestion, the general mechanisms for detecting congestion, then describes the types of information needs to be carried in the congestion notification message. It also analyzed the options of transmitting congestion information, either by extending existing protocols or introducing new protocols.

[I-D.liu-rtgwg-path-aware-remote-protection] describes the framework of path-aware remote protection. It contains the routing plane, the forwarding plane and the remote failure notification. Similar to [I-D.cheng-rtgwg-adaptive-routing-framework], path awareness is required in the routing plane and forwarding plane for rapid switchover. It gives the requirements on remote link detection that the failure notification should be independent of routing protocols, and broadcast flooding should be avoided. It also talks about the protection scope of remote protection, which may have impacts on the speed and propagation of failure notification.

[I-D.li-rtgwg-distributed-lossless-framework] analyzes the challenges in building ultra large scale data centers for AI training, and introduces the scenarios of distributed AIDC networks. Then it proposes a framework and a set of key technologies for building lossless and reliable interconnection between multiple data centers. Global load balancing, precise flow-control and packet loss detection are mentioned as key mechanisms.

This shows that the scope of the framework documents are different, while some of the content are overlapped. There is a possibility to combine the existing framework documents to build a complete framework which includes both congestion and protection, and covers both intra-DC and inter-DC scenarios.

2.3. Information Model

[I-D.zhou-rtgwg-perceptive-routing-information] defines the information model for perceptive routing (PR), which provides the necessary information and relationship of the components in the implementation of adaptive routing systems. It offers a common information model for representing the state of the network, allowing devices to communicate critical information such as failures, congestion, and optimal paths, facilitating dynamic and automated decision-making. The information model for a PR sensing node includes a set of local information and network-level information which can be used to evaluate whether a PR notification needs to be generated and sent. The information model for a PR routing node includes a set of decisions and behaviors to be made by a PR routing node on receipt of the PR notification.

2.4. Solutions

In the solution space, there are several documents which propose the mechanisms for routing in AI networks include topology-specific mechanisms, extensions to routing protocols and the new protocols for the notification of network status.

2.4.1. Topology-specific Routing Mechanisms

[I-D.agt-rtgwg-dragonfly-routing] provides an overview of the Dragonfly+ topology, and describes the routing and forwarding mechanisms in the Dragonfly+ topology, which relies heavily on non-minimal routing and adaptive load balancing for efficient use of available network capacity. It uses existing routing mechanisms such as VRF, route leaking and EBGp to achieve route propagation control and routing policy. In terms of adaptive load balancing, the purpose is to fill paths starting from high priority, and try to move flows from congested paths as a reaction to congestion. It requires that

adaptive load balancing be able to work without complete knowledge of network link utilization and queue state. It also considers that adaptive routing can work as a complementary failure handling mechanism faster than routing convergence. While the detailed adaptive routing and load balancing mechanisms are left to other documents.

2.4.2. Extensions to Routing Protocols

[I-D.xu-idr-fare] proposes extensions to BGP to carry end-to-end path bandwidth within the data center fabric for adaptive routing. In the draft a new type of BGP Extended Community is defined, and its usage in BGP route update distribution is specified using examples of 3-stage and 5-stage Clos networks. With the information of path bandwidth and link bandwidth, weighted ECMP load balancing can be performed.

[I-D.wang-idr-next-next-hop-nodes] proposes extensions to BGP to carry the next-next hop nodes associated with a given BGP next hop. One usage of the next-next hops information is for global load balancing (GLB) in a Clos network, where load balancing based on local next-hop information cannot mitigate the congestion, and it requires help from the previous hop(s) to shift the traffic to alternative next-hop nodes towards a next-next hop node. The next-next hop information is encoded as a new characteristic code of the BGP Next Hop Dependent Characteristics Attribute.

2.4.3. New Protocols for Fast Notification

[I-D.wh-rtgwg-adaptive-routing-arn] specifies Adaptive Routing Notification (ARN) as a general mechanism to proactively disseminate congestion/failure detection and elimination information for remote nodes to perform re-routing policies. An ARN message contains two kinds of information: information reflecting the type of notification (congestion or failure) and quantifiable metrics (e.g., congestion level), and information carrying details about the affected object (e.g., affected traffic, affected paths). The ARN messages can be sent using unicast or multicast to other network nodes. The format of the ARN packets and its processing on the sending and receiving nodes are also specified. The impact to route oscillation and packet reordering caused by ARN are for further study.

[I-D.liu-rtgwg-adaptive-routing-notification] describes the information carried in Adaptive Routing Notification (ARN) messages and the mechanisms of delivering ARN message in the network. The draft gives three options, each of which specifies the information carried in the ARN message and the mechanism of sending the message to specific network nodes. The complexity and overhead in

implementation are also analyzed. It also introduces an ARN TAG mechanism to control the enabling of ARN mechanism on specific traffic flows.

[I-D.zzhang-rtgwg-router-info] specifies a generic mechanism for a router to advertise some information to its neighbors. One use case is to advertise link or path information to allow receiving node to better react to network changes. The draft firstly analyzes the requirements for the information advertisement, then chooses to use UDP as a better choice comparing to IGP. The format of the message and the contained information are defined in the draft. How the IP address of the target nodes are obtained, and the processing on the receiving nodes are considered out of scope of the draft.

3. Implementations in Industry

One of the most prominent applications of fast notification is adaptive routing, which has recently gained significant traction in Ethernet-based Artificial Intelligence Data Centers (AIDCs). These data centers require real-time network information to dynamically handle the unpredictable and bursty traffic of AI/ML applications. The following sections highlight some notable implementations of adaptive routing in modern data center environments.

3.1. DLB and GLB

Dynamic Load Balancing (DLB) is a mechanism that selects the next hop for packets based on the quality of the local switch port or other local information. Global Load Balancing (GLB) extends this approach by considering the quality of downstream paths when selecting the next hop, thereby optimizing traffic distribution and improving overall network efficiency. The DLB and GLB mechanisms are implemented by many data center switches, including those from Broadcom [GLB-Broadcom], Juniper [GLB-Juniper], and Nvidia [GLB-NVIDIA].

3.2. VRF-based Adaptive Routing

Huawei's CloudEngine series switches implement adaptive routing through a VRF-based architecture [VRF-AR]. This design maintains three distinct routing tables on each device: one for shortest paths, one for non-shortest paths, and a combined table for both. Path selection is dynamically adjusted based on real-time network conditions, including both the local port status and global congestion status. The latter is communicated via Adaptive Routing Notifications (ARN), allowing for intelligent, congestion-aware routing decisions that enhance overall network performance and resiliency.

3.3. CONGA

Cisco developed a solution [CONGA] for datacenter fabrics. CONGA is a network-based, distributed, congestion-aware load balancing mechanism designed for datacenter Clos topologies and network virtualization overlays. It is CONGA splits TCP flows into flowlets, estimates real-time congestion on fabric paths using feedback from remote switches, and dynamically allocates flowlets to optimal paths.

3.4. Centralized TE and E-ECMP

Meta has developed several solutions such as centralized Traffic Engineering (TE) and Enhanced ECMP (E-ECMP) which are specifically designed for AI workloads [TE-EECMP].

In the centralized TE approach, real-time workload and network topology information are collected and transmitted to the control plane. The TE engine then executes the Constrained Shortest Path First (CSPF) algorithm to generate optimized flow placements every 30 seconds. The resulting flow placement policy overrides the default BGP routes on each switch, with BGP routing decisions serving exclusively as a backup mechanism.

E-ECMP is designed to address the low entropy inherent in AI workload flows. To achieve this, switches are configured to additionally hash the QP field of RoCE packets. Furthermore, NIC-to-NIC flows are divided into multiple flows to increase the number of QPs, thereby enhancing load distribution.

4. Summary and Potential Work

The analysis about the current state of the art for routing in AI networks shows that "Adaptive Routing" is a vague term and has different meanings in different documents or implementations. In some cases, it refers to dynamic load balancing taking the link congestion status into consideration. While in some other cases, it refers to fast switchover due to network failure. As claimed in some documents, adaptive routing is faster than route convergence, the functionalities specified in the documents are not directly related to routing or path computation. In the industry, global load balancing (GLB) is used in many solutions, while it does not cover the failure cases. It seems that a better term may be needed in IETF to more accurately reflect the functionality.

According to the framework and solutions documents, it seems the related work mainly includes: routing extensions for more visibility in network topology and capacity information, fast notification of network congestion or failure conditions, and dynamic traffic

engineering and load balancing mechanisms. In some gap analysis and problem statements, congestion control is also considered as one of the problems to be solved. While since congestion management belongs to the WIT area in IETF, it is not clear whether it can be pursued together with other functions in the RTG area.

In many of the analyzed documents, it is assumed that the underlay routing is based on EBGp, and extensions to BGP for the advertisement of additional network information are proposed. Whether other routing protocol options (e.g., IGP, IBGP, BGP-SPF, RIFT etc.) also need to be investigated is something for further consideration.

In terms of load balancing, currently most of the documents and solutions focus on the load balancing over ECMP paths, while in some topologies (such as Dragonfly and Dragonfly+), non-ECMP paths may also need to be taken into consideration.

It seems there is common interest in the fast notification mechanism for traffic engineering and load balancing. This may be something a new initiative in IETF could start with, and there is some open questions for further discussion. As mentioned in some of the documents, congestion notification is required for dynamic load balancing or flow redirect, and failure notification is required for fast switchover. Currently it is not clear whether it is possible to provide a general mechanism for the notification of both the congestion and failure conditions, or there are enough differences between the two cases that separate mechanisms are needed. Moreover, further investigation is needed on whether a new protocol is needed for fast notification, or extensions based on existing protocols would also meet some of the requirements.

5. Security Considerations

TBD

6. IANA Considerations

There are no requested IANA actions.

7. Acknowledgments

The authors would like to thank Xuesong Geng and Hang Shi for their review and discussion of this document.

8. Informative References

[GLB-Broadcom]

"Cognitive routing in the Tomahawk 5 data center switch", n.d., <<https://www.broadcom.com/blog/cognitive-routing-in-the-tomahawk-5-data-center-switch>>.

[GLB-Juniper]

"Global Load Balancing (GLB)", n.d., <<https://www.juniper.net/documentation/us/en/software/junos/ai-ml-evo/topics/topic-map/glb.html>>.

[GLB-NVIDIA]

"Turbocharging Generative AI Workloads with NVIDIA Spectrum-X Networking Platform", n.d., <<https://developer.nvidia.com/blog/turbocharging-ai-workloads-with-nvidia-spectrum-x-networking-platform>>.

[VRF-AR]

"What Is Dragonfly Adaptive Routing?", n.d., <<https://info.support.huawei.com/info-finder/encyclopedia/en/Dragonfly+Adaptive+Routing.html>>.

[CONGA]

"CONGA-Distributed Congestion-Aware Load Balancing for Datacenters", n.d., <<https://dl.acm.org/doi/pdf/10.1145/2740070.2626316>>.

[TE-EECMP]

"RDMA over Ethernet for Distributed Training at Meta Scale", n.d., <<https://dl.acm.org/doi/10.1145/3651890.3672233>>.

[I-D.hcl-rtgwg-ai-network-problem]

Huo, P., Chen, G., Lin, C., and Z. Jiang, "Gap Analysis, Problem Statement, and Requirements in AI Networks", Work in Progress, Internet-Draft, draft-hcl-rtgwg-ai-network-problem-01, 23 August 2024, <<https://datatracker.ietf.org/doc/html/draft-hcl-rtgwg-ai-network-problem-01>>.

[I-D.cheng-rtgwg-ai-network-reliability-problem]

Cheng, W., Lin, C., wangwenxuan, and B. Xu, "Reliability in AI Networks Gap Analysis, Problem Statement, and Requirements", Work in Progress, Internet-Draft, draft-cheng-rtgwg-ai-network-reliability-problem-02, 3 November 2024, <<https://datatracker.ietf.org/doc/html/draft-cheng-rtgwg-ai-network-reliability-problem-02>>.

[I-D.wang-rtgwg-dragonfly-routing-problem]

Wang, R., Lin, C., wangwenxuan, and W. Cheng, "Routing mechanism in Dragonfly Networks Gap Analysis, Problem Statement, and Requirements", Work in Progress, Internet-

Draft, draft-wang-rtgwg-dragonfly-routing-problem-02, 4 September 2024, <<https://datatracker.ietf.org/doc/html/draft-wang-rtgwg-dragonfly-routing-problem-02>>.

[I-D.cheng-rtgwg-adaptive-routing-framework]

Cheng, W., Lin, C., Wang, K., Ye, J., Zhuang, R., and P. Huo, "Adaptive Routing Framework", Work in Progress, Internet-Draft, draft-cheng-rtgwg-adaptive-routing-framework-03, 20 October 2024, <<https://datatracker.ietf.org/doc/html/draft-cheng-rtgwg-adaptive-routing-framework-03>>.

[I-D.liu-rtgwg-path-aware-remote-protection]

Liu, Y., Lin, C., Chen, M., Zhang, Z., Wang, K., and Z. He, "Path-aware Remote Protection Framework", Work in Progress, Internet-Draft, draft-liu-rtgwg-path-aware-remote-protection-02, 13 September 2024, <<https://datatracker.ietf.org/doc/html/draft-liu-rtgwg-path-aware-remote-protection-02>>.

[I-D.li-rtgwg-distributed-lossless-framework]

Li, C., Ji, S., and K. Zhu, "Framework of Distributed AIDC Network", Work in Progress, Internet-Draft, draft-li-rtgwg-distributed-lossless-framework-00, 21 October 2024, <<https://datatracker.ietf.org/doc/html/draft-li-rtgwg-distributed-lossless-framework-00>>.

[I-D.zhou-rtgwg-perceptive-routing-information]

Zhou, T., Li, D., and X. Geng, "Perceptive Routing Information Model", Work in Progress, Internet-Draft, draft-zhou-rtgwg-perceptive-routing-information-00, 18 October 2024, <<https://datatracker.ietf.org/doc/html/draft-zhou-rtgwg-perceptive-routing-information-00>>.

[I-D.agt-rtgwg-dragonfly-routing]

Afanasiev, D., Roman, and J. Tantsura, "Routing in Dragonfly+ Topologies", Work in Progress, Internet-Draft, draft-agt-rtgwg-dragonfly-routing-01, 4 March 2024, <<https://datatracker.ietf.org/doc/html/draft-agt-rtgwg-dragonfly-routing-01>>.

[I-D.xu-idr-fare]

Xu, X., Hegde, S., He, Z., Wang, J., Huang, H., Zhang, Q., Wu, H., Liu, Y., Xia, Y., Wang, P., and Tiezheng, "Fully Adaptive Routing Ethernet using BGP", Work in Progress, Internet-Draft, draft-xu-idr-fare-02, 1 September 2024, <<https://datatracker.ietf.org/doc/html/draft-xu-idr-fare-02>>.

[I-D.wang-idr-next-next-hop-nodes]

Wang, K., Haas, J., Lin, C., and J. Tantsura, "BGP Next-next Hop Nodes", Work in Progress, Internet-Draft, draft-wang-idr-next-next-hop-nodes-02, 2 December 2024, <<https://datatracker.ietf.org/doc/html/draft-wang-idr-next-next-hop-nodes-02>>.

[I-D.wh-rtgwg-adaptive-routing-arn]

Wang, H., Huang, H., Geng, X., Xu, X., and Y. Xia, "Adaptive Routing Notification", Work in Progress, Internet-Draft, draft-wh-rtgwg-adaptive-routing-arn-03, 13 September 2024, <<https://datatracker.ietf.org/doc/html/draft-wh-rtgwg-adaptive-routing-arn-03>>.

[I-D.liu-rtgwg-adaptive-routing-notification]

Liu, Y., lihesong, and W. Duan, "Adaptive Routing Notification for Load-balancing", Work in Progress, Internet-Draft, draft-liu-rtgwg-adaptive-routing-notification-01, 20 October 2024, <<https://datatracker.ietf.org/doc/html/draft-liu-rtgwg-adaptive-routing-notification-01>>.

[I-D.zzhang-rtgwg-router-info]

Zhang, Z. J., Wang, K., Lin, C., and N. Vaidya, "Advertising Router Information", Work in Progress, Internet-Draft, draft-zzhang-rtgwg-router-info-01, 18 September 2024, <<https://datatracker.ietf.org/doc/html/draft-zzhang-rtgwg-router-info-01>>.

Authors' Addresses

Jie Dong
Huawei Technologies
No. 156 Beiqing Road
Beijing
China
Email: jie.dong@huawei.com

Dan Li
Tsinghua University
Beijing
China
Email: tolidan@tsinghua.edu.cn

Qinru Shi
Huawei Technologies
No. 156 Beiqing Road
Beijing
China
Email: shiqinru@huawei.com

PengFei Huo
ByteDance
Beijing
China
Email: huopengfei@bytedance.com