Zero Trust Working Group
Internet-Draft
Intended status: Informational
Expires: 9 July 2026

X. Li
A. Wang
J. Chen
W. Lin
China Telecom

5 January 2026

Consideration of Applying Zero Trust Philosophy in Network Infrastructure
draft-li-zt-consideration-01

Abstract

Network security has traditionally relied on a perimeter-centric model, assuming that traffic originating within the network can be implicitly trusted. This model is fundamentally challenged by modern, highly distributed, and software-driven network environments where internal compromise is a realistic and high-impact threat scenario. This document examines the critical limitations of edge-only network protection and the systemic risks that arise from insufficient internal validation. Once the network perimeter is bypassed, the absence of internal protection mechanisms facilitates rapid lateral movement, impersonation of network entities, and interference with critical control and management functions.

The document argues that Zero Trust (ZT) principles, which mandate continuous, dynamic verification of all entities and communications regardless of network location, are necessary to address contemporary threat models. Deploying ZT-aligned network protection mechanisms beyond the network edge is essential to build resilient, controllable, and trustworthy networks.

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 9 July 2026.

Table of Contents

1. Introduction

Traditional network security architectures in operator and enterprise environments have long been built around a perimeter-centric protection model. In this model, security mechanisms are primarily deployed at network edges— such as access networks [RFC2827], inter-domain boundaries, or gateway nodes— under the core assumption that traffic originating inside the network can be inherently trusted once it passes the perimeter. This assumption of Implicit Trust reflected earlier network environments in which infrastructures were relatively static, tightly controlled, and operational roles were clearly

separated. In such contexts, perimeter-based protection provided a reasonable balance between security and operational complexity.

Modern networks, however, have evolved into highly distributed, virtualized, and software-driven systems. Automated orchestration, programmable control planes [RFC7426], open management interfaces, and closed-loop control systems significantly expand the internal attack surface and increase the potential impact of internal failures or compromise. As a result, threats originating from within the network can no longer be treated as exceptional or out of scope. The reliance on Implicit Trust within the network creates a structural mismatch between the threat environment and deployed protection mechanisms.

This document examines the limitations of the perimeter-centric model and the necessity of applying Zero Trust principles to network protection itself. Zero Trust rejects trust based on network location and emphasizes continuous verification of entities and communications. Applying these principles within the network enables more robust containment of compromise and improved resilience of network operations.

2. Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119] .

3. Terminology

The following terms are defined in this document:

* ZTA: Zero Trust Architecture. An evolving set of cybersecurity paradigms that move defenses from static, network-based perimeters to focus on users, assets, and resources.

* Implicit Trust: The assumption that an entity (user, device, traffic flow) is trustworthy solely because of its network location (e.g., being inside the network perimeter).

* Lateral Movement: The technique used by attackers to progressively move deeper into a network from an initial point of compromise, often by exploiting Implicit Trust.

4. Current State of Network Protection

In today's operational networks, the dominant security paradigm remains perimeter-centric. Most protection mechanisms are concentrated at the network boundary, reflecting the historical assumption of Implicit Trust for internal traffic. Common practices include:

* Traffic filtering, access control, and anomaly detection primarily enforced at ingress or egress points.[RFC2827]

* Security inspection and policy enforcement focused on customer-facing interfaces and inter-domain links.

* Limited or coarse-grained security controls within the internal network, where routers, switches, virtualized network functions, and control systems are often treated as mutually trusted.

This architectural approach originated in an era when networks were relatively static and infrastructure components were physically isolated. Under such conditions, deploying strong security controls only at the boundary was often sufficient and operationally efficient. However, the shift to virtualized, cloud-native, and software-driven networks has rendered this model increasingly fragile.

5. Risks of the Perimeter-Centric Model

A security architecture that relies primarily on edge-based protection exhibits a critical weakness: once the perimeter is breached, the internal network is left largely unprotected. This creates a "hard shell, soft interior" structure, leading to systemic risks across the network planes.

5.1. Data Plane Risks: Unrestricted Lateral Movement

The core risk is the unrestricted lateral movement of an attacker who gains an initial foothold inside the network. Because internal traffic is subject to minimal verification, a compromised node can move across internal segments with limited resistance, accessing additional systems and services. Furthermore, the lack of internal validation mechanisms means that compromised nodes can easily impersonate other network elements or services, undermining trust relationships within the network. While edge-based mechanisms address external spoofing, they do not prevent a compromised internal entity from spoofing other internal entities.

5.2. Control Plane Risks: Integrity Exposure

Internal control protocols (e.g., routing, signaling) and management interfaces are often designed with the assumption of Implicit Trust. This exposure is critical because:

* Control protocols may accept unauthenticated or insufficiently verified traffic, enabling disruption or manipulation of network operations (e.g., malicious routing updates).

* In automated and intelligent networks, incorrect or malicious internal signals can trigger large-scale misconfigurations or service disruptions, as autonomous control loops amplify the original compromise.

5.3. Management Plane Risks: API and Orchestration Vulnerability

Modern networks rely heavily on open APIs, software-defined networking (SDN) controllers [RFC7426], and automated orchestration systems. These systems manage the entire network state, making them high-value targets for attackers. If an attacker gains access to the management plane through a compromised internal entity, they can leverage the Implicit Trust to execute high-impact actions, such as reconfiguring security policies, redirecting traffic, or disabling critical network functions.

A critical and often overlooked risk in this context is the lack of robust user profile modeling for management plane access. User profile modeling establishes a baseline of legitimate behaviors, permissions, and operational contexts specific to individual accounts (e.g., network administrators, automation service accounts). This baseline includes factors such as typical access times, frequently interacted API endpoints, allowed action scopes, data access patterns, and even contextual attributes like device fingerprints or network locations. When such profiling is absent or insufficient, attackers who successfully bypass initial Zero Trust authentication (e.g., through stolen credentials, token hijacking, or session spoofing) can operate under the "legitimate" token of a compromised account while performing actions that deviate drastically from the account's intended use.

For example, a network operations account with a baseline profile limited to routine configuration checks and minor adjustments could be exploited to deploy malicious orchestration workflows, delete critical network policies, or exfiltrate sensitive management data—all without triggering alerts, as the system lacks the ability to distinguish between authorized users acting within their role and attackers misusing the account. This risk is amplified by the

automation-centric nature of modern management planes: a single compromised token with unrestricted or poorly validated access can initiate cascading, automated attacks across the network infrastructure. Unlike external threats, these actions are executed under the guise of legitimate account activity, evading traditional perimeter-focused or token-only authentication controls.

The absence of user profile modeling creates a gap in the Zero Trust framework for the management plane: while authentication verifies "who" is accessing the system, it fails to validate whether the "what" (the action being performed) aligns with the account's intended identity and operational baseline. This mismatch allows attackers to abuse legitimate credentials to undermine the integrity and availability of core network management functions, even in environments where perimeter and authentication controls are otherwise robust.

6. Necessity of Zero Trust Deployment Within the Network

Zero Trust (ZT) principles address these challenges by eliminating Implicit Trust and requiring continuous, dynamic verification across the entire network. Trust is never implicit and must be continuously reassessed based on identity, context, and behavior. This approach is necessary for network protection for several reasons:

* Elimination of Trust-by-Location: Network nodes and traffic are no longer trusted solely because they originate from internal segments, forcing explicit authentication and authorization for all interactions.

* Containment of Compromise: Security enforcement at multiple internal points limits the "blast radius" of a compromised component and restricts lateral movement, transforming the network from a soft interior to a segmented, hardened structure.

* Improved Integrity of Control and Management Functions: Continuous verification helps ensure that routing, orchestration, and monitoring systems operate on trustworthy inputs, which is vital for the stability of automated network operations.

* Resilience and Compliance: ZT provides a framework for building networks that are inherently more resilient to internal threats and better aligned with modern security compliance mandates.

Applying Zero Trust to network protection implies that internal communications, forwarding behaviors, and control interactions must be subject to security enforcement similar to that applied at the perimeter. This requires the development of network mechanisms that can enforce policy based on identity and context, rather than just network address and location.

7. Conclusion

The evolution of network architectures and threat models has rendered traditional edge-only security approaches insufficient. While perimeter defenses remain necessary, they are no longer adequate on their own. A breach at the boundary can expose the internal network to rapid and wide-ranging compromise. Adopting Zero Trust principles within the network is therefore not optional, but essential. By shifting from static, perimeter-based trust to dynamic, continuous verification across all network segments, operators can build more resilient, controllable, and trustworthy networks. Zero Trust-aligned network protection transforms security from a boundary function into an intrinsic property of the network itself, better suited to the demands of modern and future network environments.

8. Security Considerations

This document is a Problem Statement and does not propose a solution. However, the deployment of Zero Trust principles within the network introduces its own set of security and operational considerations that must be addressed by any future solution. These include:

* Performance Overhead: Continuous verification and policy enforcement at multiple internal points may introduce latency and performance overhead to the data plane.

* Reliability and Availability: The Policy Decision Point (PDP) and Policy Enforcement Point (PEP) components of a ZT architecture represent critical infrastructure. Their failure could lead to network disruption or denial of service.

* Policy Complexity: Managing fine-grained, dynamic policies across a large, distributed network is complex and requires robust automation and orchestration to avoid misconfiguration and policy conflicts.

9. IANA Considerations

TBD

10. Acknowledgement

TBD

11. References

11.1. 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>.

[RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, DOI 10.17487/RFC2827, May 2000, <https://www.rfc-editor.org/info/rfc2827>.

[RFC7426] Haleplidis, E., Ed., Pentikousis, K., Ed., Denazis, S., Hadi Salim, J., Meyer, D., and O. Koufopavlou, "Software-Defined Networking (SDN): Layers and Architecture Terminology", RFC 7426, DOI 10.17487/RFC7426, January 2015, <https://www.rfc-editor.org/info/rfc7426>.

11.2. Informative References

Authors' Addresses

Xueting Li
China Telecom
Beiqijia Town, Changping District
Beijing
Beijing, 102209
China
Email: lixt2@foxmail.com

Aijun Wang
China Telecom
Beiqijia Town, Changping District
Beijing
Beijing, 102209
China
Email: wangaj3@chinatelecom.cn

Jie Chen
China Telecom
Zhongshan Avenue, Tianhe District
Guangzhou
Email: chenj135@chinatelecom.cn

Wenhao Lin
China Telecom
Zhongshan Avenue, Tianhe District
Guangzhou
Email: liwh21@chinatelecom.cn