

BYTEDANCE USES EBPF TO ENHANCE NETWORKING PERFORMANCE



OVERVIEW

Bytedance, a global technology company operating a wide range of content platforms around the world at massive scale, faced significant challenges in ensuring performance and stability across its data centers. With over a million servers running containerized applications, the company required a networking solution that could handle high throughput while maintaining stability. By leveraging eBPF technology, Bytedance successfully implemented a decentralized networking solution that improved efficiency, scalability, and performance.

CHALLENGE

As Bytedance scaled its operations, its existing container networking solution, which relied on virtual Ethernet devices, began showing limitations. Key challenges included:

- **Performance Bottlenecks:** The virtual Ethernet-based solution introduced soft-interrupt bottlenecks in the network stack, causing inefficiencies.
- **Stability Concerns:** Ensuring stability at such a large scale was critical, but unproven solutions posed risks in production environments.
- **Kernel Version Constraints:** Upgrading to the required kernel version (6.8) was not immediately feasible due to operational constraints.

To address these challenges, Bytedance required a robust, scalable, and high-performance solution that could be deployed incrementally across its data centers.

SOLUTION

Bytedance turned to eBPF to address these challenges. eBPF, a powerful technology for dynamically and safely reprogramming the Linux kernel, enabled the company to redesign its container networking stack. Key steps in the implementation included:

- **Introducing netkit:**
 - Bytedance adopted [netkit](#), an eBPF-powered networking device introduced in version 6.8 of the Linux kernel, which provides superior performance over traditional virtual Ethernet in container networking.
 - To accommodate existing constraints and legacy deployments, netkit was backported to kernel version 5.15, ensuring compatibility with Bytedance's infrastructure.
- **Rolling Upgrade Strategy:**
 - A carefully orchestrated upgrade process ensured minimal disruption. Both the container networking interface and kernel upgrades were managed independently, allowing a gradual transition to netkit.
 - There was a concern over potential issues with mixed deployments, so compatibility between netkit and virtual Ethernet was ensured throughout the transition.
- **Handling Failures:**
 - Fallback mechanisms were implemented to handle scenarios where netkit or its associated eBPF programs failed, to ensure uninterrupted service.

The deployment of eBPF and netkit delivered significant improvements across Bytedance's data centers:

- **Performance Gains:**
 - Eliminated the last soft-interrupt in the network stack, leading to a 10% improvement in throughput.
 - Resolved issues with high CPU load and packet reordering caused by virtual Ethernet, enhancing overall efficiency.
- **Scalability and Stability:**
 - Successfully deployed netkit across dozens of clusters, demonstrating its reliability and readiness for broader adoption.
- **Operational Benefits:**
 - Simplified the networking stack, reducing maintenance overhead and improving system observability.

FUTURE PLANS

Bytedance plans to further explore the potential of eBPF, including:

- **Hardware Offloading:** Combining eBPF with hardware offloading to achieve even greater network performance.
- **Broader Use Cases:** Expanding eBPF's application beyond container networking to other areas of system optimization.

CONCLUSION

Bytedance's adoption of eBPF and netkit highlights the transformative potential of this technology in addressing large-scale networking challenges. By redesigning its networking stack and leveraging the flexibility of eBPF, Bytedance achieved measurable performance improvements while ensuring stability and scalability. This deployment serves as a testament to eBPF's capability to drive innovation in modern data center operations.

To learn more about their use of netkit, check out the [talk from eBPF Summit](#).

