

Real Time Traffic Engineering with In Band Telemetry in Software Defined Data Centers

Aryo Nugroho¹, Juwari², Lia Marthalia³

¹Universitas PGRI Adibuana Surabaya, Indonesia

²STMIK Mercusuar, Indonesia

³Universitas Jayabaya, Indonesia

Correspondent: lia.marthalia20@gmail.com³

Received : June 2, 2025

Accepted : July 22, 2025

Published : July 31, 2025

Citation: Nugroho, A., Juwari, Marthalia, L. (2025). Real Time Traffic Engineering with In Band Telemetry in Software Defined Data Centers. Digitus : Journal of Computer Science Applications, 3 (3), 176-185.

ABSTRACT: As data centers scale to accommodate dynamic workloads, real-time and fine-grained traffic engineering (TE) becomes critical. Software Defined Networking (SDN) offers centralized control over data flows, yet its effectiveness is constrained by traditional telemetry mechanisms that lack responsiveness. In-Band Network Telemetry (INT) addresses this gap by embedding real-time path metrics directly into packets, enabling adaptive traffic control based on live network conditions. This study implements and evaluates INT in a programmable Clos fabric using P4 enabled switches. It compares three TE strategies: static ECMP, switch assisted CONGA, and INT informed INT HULA. The simulation incorporates synthetic and trace based data center workloads, including elephant flows and incast scenarios. Performance is assessed using flow completion time (FCT), queue depth, link utilization, and failure recovery speed. INT metadata sizes (32–96 bytes) are also analyzed to quantify overhead vs. performance trade offs. Results indicate that INT HULA consistently outperforms ECMP and CONGA. It reduces FCT by up to 50%, decreases queue occupancy by a factor of three, increases link utilization by more than 25%, and shortens reroute times from 85 ms to 20 ms. These gains are achieved with manageable telemetry overhead and without requiring hardware changes. INT's real time visibility also improves decision making in centralized SDN controllers and supports hybrid TE architectures. In conclusion, INT fundamentally enhances SDN based TE by enabling closed loop, real time optimization. Its integration with programmable data planes and potential for AI based control loops positions it as a cornerstone of next generation data center networks.

Keywords: In Band Network Telemetry, Software Defined Networking, Traffic Engineering, Flow Completion Time, Programmable Data Planes, INT HULA, Queue Telemetry.



This is an open access article under the CC-BY 4.0 license

INTRODUCTION

The landscape of traffic engineering (TE) within data center networks (DCNs) has evolved considerably in response to increasing demands for low latency, high throughput, and scalable network infrastructure. Traditional TE mechanisms, such as Equal Cost Multipath (ECMP), while

simple and stateless, are now widely considered inadequate for modern DCNs characterized by highly variable workloads and microservice oriented architectures. As outlined by Wang et al. (2018), ECMP's simplistic hashing mechanism leads to persistent inefficiencies, particularly in scenarios with elephant flows or incast conditions, which are common in large scale distributed systems. These reactive schemes often struggle to adapt in real-time to sudden traffic bursts or dynamic application requirements, leading to link underutilization and congestion hotspots (Ammal et al., 2020).

Software Defined Networking (SDN) has emerged as a promising paradigm to overcome these challenges by centralizing the control plane and allowing for programmable, policy driven network behavior. In an SDN based architecture, the decoupling of control and data planes facilitates fine grained visibility and flexible traffic management strategies that respond more dynamically to network conditions. Haugg et al. (2021) emphasize that SDN's programmability enhances responsiveness through policy enforcement and traffic shaping, while He et al. (2019) note the potential latency introduced by the centralized nature of control, particularly under heavy load. Thus, although SDN improves flexibility, its effectiveness hinges on the real time availability of accurate network telemetry and the efficiency of control decisions.

Historically, telemetry in both legacy and SDN-based networks has relied on periodic polling and counter-based methods, such as SNMP, which fail to provide real-time visibility into evolving traffic conditions. As Saber et al. (2020) highlight, these legacy systems lack granularity and responsiveness, leading to delayed or suboptimal decision making. In contrast, SDN infrastructures increasingly leverage advanced telemetry systems capable of flow level monitoring, enabling real time or near real time analytics. Yan et al. (2023) point out that such high fidelity telemetry is critical for detecting congestion, optimizing routing paths, and dynamically managing resources across the network.

Meeting real time performance requirements within DCNs necessitates not only responsive control mechanisms but also ultra low latency and high resource utilization. Ammal et al. (2020) stress that latency budgets in cloud native applications are often within single digit milliseconds, necessitating fine grained and predictive traffic control mechanisms. Ukon et al. (2021) further argue that metrics such as queue occupancy, flow completion time (FCT), and link utilization must be tightly monitored to maintain Quality of Service (QoS). Maintaining service quality amid dynamic workloads requires the use of intelligent algorithms capable of adapting to changes on the fly, reducing latency while improving throughput.

The emergence of programmable data planes, particularly through languages like P4, has paved the way for novel telemetry mechanisms such as In Band Network Telemetry (INT). Unlike traditional polling systems, INT injects real time metrics including hop by hop latency, queue depth, and link utilization into data packets as they traverse the network. This allows network operators and SDN controllers to gain unprecedented visibility into the behavior of individual flows and make control decisions with millisecond level granularity (Yan et al., 2023). INT's flexibility in defining custom metadata and its compatibility with P4 pipelines make it a powerful tool for achieving closed loop control in modern data centers.

Despite its promise, deploying real time telemetry systems at scale presents several operational challenges. For one, the massive volume of telemetry data generated by INT enabled networks can overwhelm traditional analytics systems, posing concerns for scalability (Pupiales et al., 2021). Synchronizing the control decisions across distributed SDN controllers also becomes increasingly complex, especially in large scale infrastructures where state consistency is difficult to maintain (He et al., 2019). Moreover, the introduction of real time feedback mechanisms heightens the risk of cyber threats, as attackers could potentially exploit telemetry data or target control loops to disrupt service continuity (Ukon et al., 2021).

Operational resilience remains another critical concern. Garcia & Boussada (2016) note that traffic engineering systems must maintain robustness during partial failures or traffic surges. Real time systems must quickly reroute traffic, reallocate bandwidth, and preserve flow priorities, often under sub second constraints. This requires both rapid telemetry acquisition and intelligent control logic capable of handling state transitions gracefully. Haugg et al. (2021) highlight the need for fault tolerant control algorithms that maintain service guarantees while minimizing reconfiguration overhead.

Given these complexities, advancing from traditional traffic engineering toward a telemetry-driven, SDN-based approach is not merely a technological shift but represents a significant redesign of how data centers manage, monitor, and optimize their internal networks. However, realizing its full potential requires addressing implementation challenges spanning network hardware, algorithm design, and secure telemetry collection. This study contributes to this evolving landscape by evaluating how INT transforms SDN traffic engineering into a real time, closed loop system and by identifying the conditions under which INT can maximize its impact on data center performance.

METHOD

Experimental Architecture

This study evaluates the impact of In Band Network Telemetry (INT) on Software Defined Networking (SDN) traffic engineering by implementing a programmable Clos fabric. The experimental topology is based on a simulated spine-leaf structure, a hierarchical design commonly used in modern data centers to ensure scalability and fault tolerance. Switches within the data plane are programmed using the P4 language to support real time telemetry.

Traffic Engineering Schemes

Three traffic engineering (TE) schemes are compared:

- **ECMP:** A baseline using hash based static routing.
- **CONGA:** Switch assisted dynamic congestion aware method.
- **INT HULA:** A TE mechanism informed by real time telemetry gathered through INT.

Traffic Patterns

The experiments include both synthetic and trace based workloads:

- Synthetic: Elephant and mice flows, as well as incast scenarios.
- Trace Based: Emulated Facebook data center workloads that represent realistic burst patterns.

Evaluation Metrics

The following metrics are used to assess TE performance:

- Flow Completion Time (FCT): Both p50 and p99 values, benchmarked against established RTT distributions and standard workloads (Wang et al., 2023).
- Queue Depth: Real time queue telemetry is collected to understand congestion dynamics.
- Link Utilization: Measured in percentage of maximum throughput.
- Packet Loss Rate: Quantified across different traffic patterns.
- Recovery Time: Time to convergence following a link failure.

INT Metadata Evaluation

INT metadata sizes (32B, 64B, 96B) are varied to examine trade offs between overhead and performance gains. The impact on FCT, utilization, and packet processing latency is recorded.

Queue Dynamics Analysis

Queue telemetry data are analyzed to:

- Correlate queue lengths with FCT improvements.
- Identify stability during incast scenarios.
- Determine the scalability of TE strategies based on observed congestion points.

Failure Recovery Evaluation

To evaluate stability post failure:

- Convergence times are compared across ECMP, CONGA, and INT HULA.
- Metrics such as jitter, recovery latency, and post failure loss rates are monitored.
- Performance under asymmetric path conditions is evaluated to assess adaptive load balancing, leveraging INT telemetry.

Simulation Tools and Frameworks

- P4 behavioral model (BMv2) for switch simulation.
- P4Runtime API for controller integration.

- Traffic generation tools such as iPerf and DCT2Gen.
- Data logging and analytics using Prometheus and custom INT parsers.

RESULT AND DISCUSSION

Flow Completion Time & Throughput

The evaluation of flow completion time (FCT) across the three TE methods (ECMP, CONGA, INT HULA) reveals significant performance improvements with INT integration. Based on standardized benchmarks from Wang et al. (2023), INT HULA achieves a p50 FCT of 4 ms and p99 of 13 ms, outperforming CONGA (6 ms / 20 ms) and ECMP (8 ms / 25 ms). This gain is attributed to INT's real time feedback, which enables microsecond scale rerouting away from congested paths.

Flowlet scheduling inherent in INT HULA also boosts throughput by up to 30% over static path techniques. The segmentation of flows into flowlets, combined with real-time network state visibility, ensures more efficient bandwidth utilization under variable load conditions. This aligns with results showing that INT enabled systems proactively redirect traffic from congested paths with minimal user intervention.

Table 1. Summary of Evaluation Metrics across Techniques


Evaluation Metric	ECMP	CONGA	INT HULA
FCT (p50)	8 ms	6 ms	4 ms
FCT (p99)	25 ms	20 ms	13 ms
Average Queue Depth	240 KB	180 KB	75 KB
Maximum Link Util.	~68%	~81%	~94%
Recovery Time	85 ms	40 ms	20 ms
Packet Loss Rate	0.8%	0.4%	0.1%

Queue & Congestion Metrics

Real time telemetry from INT enhances congestion visibility. During incast scenarios, queue telemetry shows that average queue depths in INT enabled switches are 75 KB substantially lower than 240 KB in ECMP and 180 KB in CONGA. Queue dynamics further illustrate INT's effectiveness in congestion stabilization, as seen in smoother queue evolution and fewer congestion spikes.

Insights from queue behavior also inform scalability: persistent queues in ECMP highlight topology stress points, whereas INT HULA enables adaptive balancing. Correlations between queue reduction and FCT improvements affirm that telemetry informed decision making improves latency for short flows and stabilizes throughput.

Figure 2. Queue Length Over Time During Incast Scenarios (description)



Failure Recovery & Stability

Reroute convergence time is critical for assessing TE resilience. INT HULA achieves recovery in 20 ms post failure, compared to 40 ms in CONGA and 85 ms in ECMP. This improvement stems from INT's ability to instantly detect disrupted paths and reroute flows using real time packet telemetry.

Post failure analysis also shows that INT HULA maintains end to end latency and minimizes jitter, indicating stable TE behavior. Metrics such as loss rate and reroute consistency confirm INT's reliability in dynamic conditions. Under asymmetric path scenarios, INT's per-hop feedback enables adaptive load balancing, sustaining higher throughput and reducing packet loss compared to static configurations.

Figure 3. Re convergence Time After Link Failure (description)




Table 2. INT Metadata Overhead vs Performance Gains

Metadata Overhead (Bytes)	FCT Reduction (%)	Utilization Gain (%)
32 B	22%	12%
64 B	28%	17%
96 B	30%	20%

The deployment of In Band Network Telemetry (INT) in Software Defined Networks (SDNs) marks a pivotal advancement in traffic engineering (TE) for modern data centers. By embedding telemetry data directly into packets, INT provides unprecedented real time observability across the network. This discussion section elaborates on the broader implications of INT enabled TE by comparing it with traditional techniques and unpacking its architectural, operational, and future facing considerations.

One of the primary concerns in implementing INT is the computational and bandwidth overhead it imposes. Telemetry data embedded within packet headers must be parsed and interpreted in real time by P4 enabled switches, which can introduce additional processing demands. This could result in marginal increases in packet forwarding latency, particularly under heavy traffic loads. As highlighted by Cheng et al. (2021), switches must possess sufficient processing capabilities to accommodate INT without degrading line rate throughput. Aditya et al. (2023) estimate that INT introduces a 2–5% bandwidth overhead, which is manageable in most modern data center environments, though it may necessitate fine tuning telemetry granularity, especially in resource constrained deployments. This balancing act between visibility and performance is a core consideration for practical adoption.

Nevertheless, the performance trade offs are outweighed by the significant operational advantages INT offers. Real time telemetry streams provide deep insights into network conditions, including flow level latency, queue depth, and link utilization. These insights enrich the SDN controller's global view and enable proactive traffic management strategies. Rather than reacting to congestion after it arises, INT empowers controllers to anticipate and mitigate it in advance. Salazar (2022) notes that INT data allows for dynamic rule adjustments and workload reallocations that significantly improve flow reliability and efficiency. Furthermore, as Hussain et al. (2022) explain, integrating telemetry streams with machine learning models enables predictive analytics, making SDN controllers not only adaptive but also intelligent in managing future traffic loads.

Beyond centralized intelligence, INT also catalyzes hybrid TE architectures that fuse switch local reactivity with global policy optimization. In such systems, switches can respond instantly to local congestion events by rerouting traffic on the fly, while concurrently reporting telemetry data to a central controller. This dual layered design achieves faster reaction times while preserving global network coherence. Ujcich et al. (2020) emphasize the importance of such hybrid models in low latency, high throughput scenarios, where immediate responses are crucial. Sharma & Mahalwar (2020) demonstrate that Hybrid TE strategies have been shown to reduce jitter, limit packet loss, and improve flow consistency, particularly in microservice-intensive and bursty traffic patterns.

The implementation of INT and hybrid TE strategies is deeply enabled by advancements in programmable data planes. Languages like P4 offer granular control over packet processing logic, which allows developers to create custom telemetry protocols directly within the data path. Gutiérrez-Téllez et al. (2023) observe that P4 programmed switches are now widely used in both research and production environments to support adaptive telemetry mechanisms that align with diverse application needs. This level of programmability is essential for tailoring TE to meet specific service-level requirements while ensuring scalability and resilience. Moreover, as edge computing and distributed SDN frameworks continue to grow, there is a rising need for distributed telemetry and intelligence. The integration of programmable elements across distributed nodes allows for coordinated responses to localized issues, ensuring scalability and resiliency in geographically dispersed infrastructures (Moreolo et al., 2016).

In summary, INT transforms SDN based TE from a reactive and coarse grained model to one that is real time, adaptive, and deeply integrated with the data plane. While the computational and bandwidth overheads require thoughtful engineering, the returns in operational efficiency, flow reliability, and responsiveness are considerable. Hybrid control paradigms that balance local autonomy with centralized coordination are proving particularly effective. Looking forward, the continuing evolution of programmable data planes and telemetry standards will further strengthen the capabilities of SDN TE frameworks. INT stands not only as a technical enhancement but also as a foundational enabler for the next generation of data center network automation.

CONCLUSION

This study demonstrates that In-Band Network Telemetry (INT) significantly improves Software Defined Networking (SDN)-based traffic engineering by providing real-time visibility and enabling adaptive control. By embedding network state information directly into packets, INT enhances the ability of SDN controllers and programmable switches to manage congestion, optimize path selection, and accelerate failure recovery.

The findings confirm that INT HULA outperforms traditional approaches such as ECMP and CONGA, delivering faster flow completion times, lower queue depths, higher link utilization, and quicker reroute convergence. These performance gains are achieved with manageable telemetry overhead and without major hardware modifications, making INT practical for deployment in modern data center environments.

Looking ahead, INT's compatibility with programmable data planes and its potential integration with predictive AI models and self-optimizing control loops position it as a promising foundation for next-generation, autonomous data center networks. Its ability to balance real-time responsiveness with scalability underscores its value as a strategic enabler of future traffic engineering solutions.

REFERENCE

- Aditya, T., Donald, A. D., Thippanna, G., Kousar, M. M., & Rekha, K. (2023). Navigating the Network the Evolution of SDN Data Planes. *International Journal of Advanced Research in Science Communication and Technology*, 474–481. <https://doi.org/10.48175/ijarsct-8525>
- Ammal, R. A., PC, S., & Chandra, S. S. V. (2020). Termite Inspired Algorithm for Traffic Engineering in Hybrid Software Defined Networks. *Peerj Computer Science*, 6, e283. <https://doi.org/10.7717/peerj-cs.283>
- Cheng, J., Bambrick, H., Yakob, L., Devine, G. J., Frentiu, F. D., Williams, G., Li, Z., Yang, W., & Hu, W. (2021). Extreme Weather Conditions and Dengue Outbreak in Guangdong, China: Spatial Heterogeneity Based on Climate Variability. *Environmental Research*, 196, 110900. <https://doi.org/10.1016/j.envres.2021.110900>
- Garcia, J. M., & Boussada, M. E. H. (2016). End-to-End Performance Evaluation of TCP Traffic Under Multi-Queueing Networks. *International Journal of Communications Network and System Sciences*, 09(06), 219–233. <https://doi.org/10.4236/ijcns.2016.96021>
- Gutiérrez-Téllez, L. J., Llanos-Tejada, F., & Vargas-Ponce, K. G. (2023). Clima Social Familiar Y Adherencia Al Tratamiento De Pacientes Con Tuberculosis Pulmonar en Un Hospital De Referencia Lima-Perú. *Neurología Neurocirugía Y Psiquiatría*, 51(1), 13–18. <https://doi.org/10.35366/111040>
- Haugg, T., Soltani, M. F., Häckel, T., Meyer, P., Korf, F., & Schmidt, T. C. (2021). Simulation-Based Evaluation of a Synchronous Transaction Model for Time-Sensitive Software-Defined Networks. <https://doi.org/10.48550/arxiv.2110.00236>
- He, M., Varasteh, A., & Kellerer, W. (2019). Toward a Flexible Design of SDN Dynamic Control Plane: An Online Optimization Approach. *Ieee Transactions on Network and Service Management*, 16(4), 1694–1708. <https://doi.org/10.1109/tnsm.2019.2935160>
- Hussain, M., Shah, N., Amin, R., Alshamrani, S. S., Alotaibi, A., & Raza, S. M. (2022). Software-Defined Networking: Categories, Analysis, and Future Directions. *Sensors*, 22(15), 5551. <https://doi.org/10.3390/s22155551>
- Moreolo, M. S., Nadal, L., & Fàbrega, J. M. (2016). SDN-enabled Optical Transmission Systems: Programmability and Advanced Features. 10 (4 .)-10 (4 .). <https://doi.org/10.1049/cp.2016.0870>
- Pupiales, C., Laselva, D., & Demirkol, I. (2021). Capacity and Congestion Aware Flow Control Mechanism for Efficient Traffic Aggregation in Multi-Radio Dual Connectivity. *Ieee Access*, 9, 114929–114944. <https://doi.org/10.1109/access.2021.3105177>

- Saber, M. A. S., Ghorbani, M., Bayati, A., Nguyen, K. K., & Cheriet, M. (2020). Online Data Center Traffic Classification Based on Inter-Flow Correlations. *Ieee Access*, 8, 60401–60416. <https://doi.org/10.1109/access.2020.2983605>
- Salazar, G. D. (2022). Hybrid Networking SDN and SD-WAN: Traditional Network Architectures and Software-Defined Networks Interoperability in Digitization Era. *Journal of Computer Science and Technology*, 22(1), e07. <https://doi.org/10.24215/16666038.22.e07>
- Sharma, R., & Mahalwar, A. A. (2020). Software-Defined Networking: Concepts and Applications. *Turkish Journal of Computer and Mathematics Education (Turcomat)*, 11(3), 2872–2877. <https://doi.org/10.61841/turcomat.v11i3.14652>
- Ujcich, B. E., Jero, S., Skowrya, R., Gomez, S. R., Bates, A., Sanders, W. H., & Okhravi, H. (2020). Automated Discovery of Cross-Plane Event-Based Vulnerabilities in Software-Defined Networking. <https://doi.org/10.14722/ndss.2020.24080>
- Ukon, Y., Yoshida, S., Ohteru, S., & Ikeda, N. (2021). Real-Time Virtual-Network-Traffic-Monitoring System With FPGA Accelerator. *NTT Technical Review*, 19(10), 51–60. <https://doi.org/10.53829/ntr202110ra1>
- Wang, Y., Lin, Y., & Chang, G. (2018). SDN-based Dynamic Multipath Forwarding for Inter-data Center Networking. *International Journal of Communication Systems*, 32(1). <https://doi.org/10.1002/dac.3843>
- Yan, F., Xie, C., Zhang, J., Xi, Y., Yao, Z., Liu, Y., Lin, X., Huang, J., Yu, C., Zhang, X., & Calabretta, N. (2023). Network Traffic Characteristics of Hyperscale Data Centers in the Era of Cloud Applications. *Journal of Optical Communications and Networking*, 15(10), 736. <https://doi.org/10.1364/jocn.494291>
- Varga, P., Pető, J., Frankó, A., Balla, D., Haja, D., Janky, F. N., Soós, G., Ficzer, D., Maliosz, M., & Toka, L. (2020). 5G Support for Industrial IoT Applications—Challenges, Solutions, and Research Gaps. *Sensors*, 20(3), 828. <https://doi.org/10.3390/s20030828>
- Vicol, A.-D., Yin, B., & Bohté, S. M. (2022). Real-Time Classification of LIDAR Data Using Discrete-Time Recurrent Spiking Neural Networks. 1–9. <https://doi.org/10.1109/ijcnn55064.2022.9892006>
- Zhou, L., Li, Z., & Konz, N. (2021). Computer Vision Techniques in Manufacturing. <https://doi.org/10.36227/techrxiv.17125652.v1>