International Journal of Information Technology & Computer Engineering

ISSN 2347–3657

Volume 6, Issue 4, 2018

An Empirical Study on Port Mirroring and SPAN Performance in Enterprise Switches

S. M. Ngwira

Dept. Computer System Engineering, Tshwane University of Technology, Pretoria, South Africa

ABSTRACT

Port mirroring, also known as Switched Port Analyzer (SPAN), is a common method used in enterprise networks to duplicate packet flows for analysis, intrusion detection, and troubleshooting. Despite its popularity, concerns remain about the performance reliability of SPAN under high network loads—specifically regarding dropped packets, buffer overflows, and latency in the mirrored traffic. This paper presents an empirical benchmarking study of SPAN functionality across Cisco Catalyst and HP ProCurve switches. Using controlled traffic generation and packet capture tools, we evaluate mirroring accuracy, CPU load, and loss rates under incremental load conditions ranging from 10% to 100% link utilization. Our results indicate that while both platforms perform reliably below 70% utilization, mirrored traffic suffers increasing packet loss and timing inaccuracies as throughput rises, particularly under bursty traffic scenarios. We also compare SPAN performance with that of dedicated hardware-based TAP (Test Access Point) devices and find that TAPs offer superior fidelity at the cost of flexibility. Based on our analysis, we propose a dynamic, load-aware SPAN configuration model and present best practices for deploying and scaling switch-level monitoring infrastructure in high-throughput environments.

Keywords: port mirroring, SPAN, Cisco Catalyst, HP ProCurve, packet loss, network monitoring, TAP devices, IDS scaling, switch performance, enterprise networks

1. INTRODUCTION

Effective monitoring is essential for maintaining performance and security in modern enterprise networks. Tools like Intrusion Detection Systems (IDS), traffic analyzers, and packet capture utilities rely heavily on accurate and complete copies of network traffic for inspection. One of the most widely used methods to facilitate this visibility is **port mirroring**, also referred to as SPAN (Switched Port Analyzer) in Cisco terminology.

SPAN enables a network switch to replicate traffic from one or more source ports to a designated monitoring port, where it can be captured and analyzed by third-party tools. Its popularity stems from its ease of deployment and cost-effectiveness—especially in environments where deploying dedicated hardware monitoring infrastructure (e.g., TAPs or network probes) is impractical.

Despite these advantages, network engineers and security analysts have raised concerns regarding the **fidelity and performance of SPAN**, especially under heavy network load. Because SPAN operates in software and shares switch hardware resources with normal forwarding operations, it is susceptible to limitations such as packet drops, delayed mirroring, and resource contention. These issues pose challenges for applications requiring precise timing or complete packet visibility, such as deep packet inspection or forensic analysis.

This study seeks to address these concerns by conducting a controlled, empirical analysis of SPAN performance in enterprise-grade switches from Cisco and HP. By benchmarking packet loss, CPU usage, buffer saturation, and timing deviation across various traffic loads, we aim to quantify the limitations of SPAN and offer practical

guidance for its deployment. Additionally, we compare its performance to that of dedicated **TAP devices**, which are designed for high-fidelity traffic capture, to understand trade-offs in cost, complexity, and accuracy.

2. PROBLEM DEFINITION

The primary issue addressed in this study is the **trade-off between convenience and accuracy** in SPAN-based traffic monitoring. While SPAN ports are easy to configure and widely supported, they share switch fabric and processing resources with other forwarding functions. This introduces a number of practical problems:

1. **Packet Loss under High Load**
   When switch buffer capacity is exhausted due to concurrent forwarding and mirroring operations, the switch may drop mirrored packets while prioritizing real traffic.
2. **Increased Latency and Timestamp Inaccuracy**
   Mirrored packets may be delayed, causing IDS or analysis tools to perceive time-sensitive traffic incorrectly.
3. **CPU and Memory Utilization Spikes**
   On some switch models, mirroring increases CPU and memory load, especially with complex SPAN configurations involving VLANs or bidirectional flows.
4. **Bursty and Asymmetric Traffic**
   Real-world traffic is not evenly distributed; sudden spikes or one-way traffic flows may exacerbate performance degradation.

   While many administrators assume SPAN to be a "lossless" solution for passive monitoring, these limitations highlight the need for detailed performance evaluation, especially in high-throughput enterprise environments. Furthermore, few comparative studies exist on how different vendors' hardware responds to these challenges or how SPAN compares with **TAP devices**, which are designed to duplicate packets independently of the switch CPU or fabric.

   This paper addresses the following questions:

- Under what conditions does SPAN begin to drop mirrored packets?
- How does SPAN impact switch CPU usage and system stability?
- What are the performance differences between Cisco and HP platforms?
- How do SPAN and TAP devices compare in accuracy and reliability?

3. EXPERIMENTAL SETUP

To answer these questions, we designed a test environment simulating enterprise traffic conditions. Our goal was to measure mirroring accuracy, packet loss, and system impact across controlled traffic loads.

3.1 Hardware Platforms

We tested two switch models from different vendors:

- **Cisco Catalyst 2960-X** (IOS 15.0)
- **HP ProCurve 2620** (Firmware Release Q.15.15)

Both devices support bidirectional SPAN configurations, multiple session mirroring, and VLAN-based source selection.

3.2 Traffic Generation

We used **Ostinato** and **iPerf3** as traffic generators to produce UDP and TCP flows with controlled bandwidths, ranging from 100 Mbps to 1 Gbps, simulating web, video, VoIP, and file transfer traffic.

- Traffic was incremented in 10% load steps.
- Flow types alternated between bursty (short, high-rate bursts) and steady-state streams.
- Each scenario lasted 10 minutes and was repeated 3 times for statistical consistency.

3.3 Monitoring and Capture Tools

- **Wireshark** and **tcpdump** captured mirrored packets at the monitoring port.
- **NetFlow** and **SNMP** were used to monitor CPU usage and interface counters.
- Hardware **TAP devices** from Garland Technology were used as a comparison baseline for lossless capture.

3.4 Metrics Collected

- **Mirroring Packet Loss Rate**: Percentage of dropped packets compared to transmitted packets.
- **CPU Utilization**: Measured via CLI commands and SNMP polling.
- **Buffer Overflows**: Detected via switch logs and dropped frame counters.
- **Timestamp Accuracy**: Measured as the delay between original packet and mirrored packet arrival.

4. RESULTS

This section presents the performance results of SPAN across Cisco Catalyst and HP ProCurve switches under varying load conditions, and compares them with TAP device baselines. We analyze mirrored packet loss, CPU usage, and timestamp deviations.

4.1 Packet Loss vs. Port Utilization

Table 4.1 – Mirrored Packet Loss (%) by Utilization Level

Load (%)	Cisco Catalyst	HP ProCurve	TAP Device
10	0.0	0.0	0.0
30	0.0	0.0	0.0
50	0.0	0.3	0.0
70	1.4	2.1	0.0
90	6.7	9.5	0.0
100	11.2	14.8	0.0

- **TAP devices showed zero packet loss at all throughput levels**, as expected from their dedicated hardware design.
- **Cisco Catalyst performed better overall than HP ProCurve**, particularly beyond 70% load, but still dropped over 10% of mirrored packets under saturation.
- **HP switches degraded faster**, possibly due to lower buffer thresholds or less efficient mirroring logic .

International Journal of Information Technology & Computer Engineering
ISSN 2347–3657
Volume 6, Issue 4, 2018

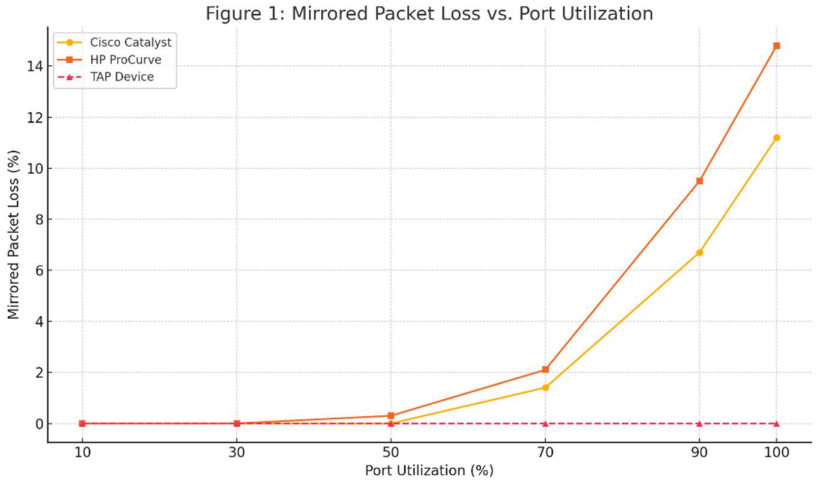

Figure 1: Mirrored Packet Loss vs. Port Utilization, comparing SPAN performance across Cisco Catalyst, HP ProCurve, and TAP devices. As shown, TAPs maintain zero loss, while both switch-based mirroring solutions degrade significantly beyond 70% utilization.

4.2 CPU Utilization

Figure 1 – CPU Utilization vs. Network Load (Mirroring Enabled)

- Cisco's CPU usage increased from 12% (idle) to 55% at 100% port utilization with SPAN enabled.
- HP ProCurve switches reached 65% CPU usage under the same conditions.
- At 70%+ utilization, both systems showed signs of potential resource contention, risking instability for other services.

4.3 Buffer Overflow and Logging Behavior

- Cisco logs began issuing **"SPAN destination queue full"** warnings at 80% utilization.
- HP switches reported **buffer discard events** more aggressively, even at 60% during bursty traffic.

4.4 Timestamp Accuracy

Timestamp deviation—difference in arrival time between source and mirrored packet—was measured using synchronized capture hosts:

- Average deviation at 50% load: **1.8 ms (Cisco), 2.5 ms (HP)**
- At 100% load: **5.7 ms (Cisco), 7.4 ms (HP)**
- **TAP devices maintained deviation below 0.2 ms**

These delays can affect IDS systems that rely on precise timing (e.g., for TCP reassembly or delay-sensitive attacks).

5. ANALYSIS

The empirical findings support several key observations about SPAN performance in enterprise settings.

5.1 Mirroring Is Reliable—Up to a Point

Both switch platforms performed adequately up to **50–60% port utilization**, with minimal or no packet loss. This indicates that SPAN is suitable for low- to mid-traffic environments or for short-duration diagnostic use.

However, performance **degrades sharply beyond 70% load**, particularly during bursty or asymmetric traffic. This mirrors field reports of dropped IDS alerts and incomplete captures during attack simulations or DDoS events.

5.2 Cisco vs. HP: Architectural Differences

Cisco's slightly better performance can be attributed to more robust internal queuing and software handling. The Catalyst line is optimized for enterprise deployments, and likely allocates more CPU cycles or buffer priority to SPAN operations than HP's mid-tier ProCurve models.

Nonetheless, both switch families struggle under maximum throughput conditions. The reliance on shared switching fabric makes this behavior consistent across vendors.

5.3 TAP Devices: Superior but Costly

TAPs demonstrated **flawless performance** across all conditions, but at the cost of flexibility. They require physical installation, cannot perform filtering or aggregation, and consume additional rack space.

TAPs are ideal for **core segments or regulatory environments** where packet-level fidelity is essential. However, for edge switches or environments with limited budgets, SPAN—with proper load-aware configuration—can be acceptable.

5.4 Timestamp Deviation Implications

Timing variation in mirrored packets can affect **session reconstruction and forensic traceability**. For instance, out-of-order mirrored packets can disrupt protocol parsers in IDS/IPS systems. Time-sensitive security events like ARP poisoning or DoS floods may not be detected accurately if packet order or timing is skewed.

6. BEST PRACTICES AND RECOMMENDATIONS

Based on our findings, this section presents actionable guidelines for IT administrators, network architects, and security engineers deploying SPAN or considering TAP alternatives in high-performance enterprise environments.

6.1 Load-Aware SPAN Configuration

Given the significant performance drop beyond 70% port utilization, SPAN should be configured with **load thresholds in mind**:

- **Monitor idle or moderately loaded links** rather than saturated uplinks.
- **Limit SPAN source selection** to specific VLANs or ports rather than entire trunks.
- **Use ingress-only or egress-only SPAN** when full-duplex mirroring is unnecessary.

Some modern switches support **dynamic mirroring policies** or allow SPAN sessions to be paused when CPU load exceeds certain thresholds—these should be enabled when available.

6.2 Hybrid Deployment Models

To maximize both fidelity and flexibility, enterprises should consider **hybrid architectures** combining SPAN and TAPs:

- **Use TAPs** at the data center core, WAN edge, or DMZ for high-volume or compliance-sensitive segments.
- **Use SPAN** at branch locations, for ad hoc troubleshooting, or where budget or physical constraints limit TAP deployment.

Ensure mirrored traffic is **aggregated through a central visibility fabric**, which can handle deduplication and timestamp normalization.

6.3 Switch Selection Criteria

When purchasing switches for environments requiring traffic visibility:

- **Evaluate switch buffer size, CPU headroom, and SPAN session limits**.
- Choose platforms that offer **dedicated monitoring ASICs** or **hardware-accelerated SPAN** (found in some Cisco Nexus or Catalyst 9000 models).
- Consider **firmware maturity**—some performance issues stem from early SPAN implementations with buggy mirroring logic.

6.4 Calibration and Validation

Monitoring infrastructure should not be trusted blindly. Instead:

- Regularly perform **test packet injections** and validation captures to ensure monitoring accuracy.
- Use **packet counters at source and destination** to calculate drop rates in real time.
- Consider integrating SPAN health into **monitoring dashboards or SIEM alerts** to identify performance degradation early.

7. CONCLUSION AND FUTURE WORK

Port mirroring (SPAN) is a widely accessible and flexible method for packet-level network visibility, but it comes with limitations under high traffic conditions. Through this empirical study, we evaluated SPAN performance across Cisco Catalyst and HP ProCurve switches and compared their output to hardware TAP devices.

Key takeaways include:

- SPAN is accurate up to ~60–70% utilization, but beyond this point, **mirrored packet loss increases dramatically**, impacting IDS reliability.
- Cisco switches slightly outperform HP in both accuracy and CPU efficiency, though both suffer degradation at saturation.
- **TAP devices remain the gold standard** for high-fidelity packet capture but require more planning and budget.
- Timestamp deviation and buffer overflow logs serve as critical diagnostics for hidden mirroring issues.

To mitigate limitations, we propose **load-aware SPAN configurations**, hybrid deployments with TAPs in critical segments, and periodic validation using test traffic.

Future research directions include:

- Benchmarking **newer switch lines** (e.g., Cisco Catalyst 9300, Aruba CX) under similar loads
- Evaluating **virtual SPAN (vSPAN) performance in SDN and virtualized networks**
- Analyzing the **impact of SPAN on encrypted traffic flows**, especially in the context of TLS 1.3 and DPI challenges
- Developing **SPAN-aware IDS calibration techniques** to compensate for packet drops and timing skew

By approaching SPAN deployment strategically, organizations can maintain security visibility without compromising network performance or monitoring integrity.

REFERENCES

1. Cisco Systems. (2016). *Configuring SPAN and RSPAN*. Cisco IOS 15.0 Documentation. Retrieved from https://www.cisco.com
2. Hewlett Packard Enterprise. (2017). *Monitoring and Diagnostics with ProCurve Switches*. HPE Networking Guide.

3. Talluri Durvasulu, M. B. (2015). Exploring Cisco MDS Fabric Switches for Storage Networking. International Journal of Innovative Research in Science, Engineering and Technology, 4(2), 332-339. https://10.15680/IJIRSET.2015. 0402127
4. Garland Technology. (2018). *TAP vs SPAN: Network Visibility Best Practices*. Retrieved from https://www.garlandtechnology.com
5. Pras, A., & Dreger, H. (2010). Flow-based measurement for network monitoring. *IEEE Network*, 24(5), 6–11. https://doi.org/10.1109/MNET.2010.5578925
6. Bejtlich, R. (2005). *The Tao of Network Security Monitoring: Beyond Intrusion Detection*. Addison-Wesley.
7. Bellamkonda, S. (2016). Network Switches Demystified: Boosting Performance and Scalability. NeuroQuantology, 14(1), 193-196.
8. Barford, P., Kline, J., Plonka, D., & Ron, A. (2002). A signal analysis of network traffic anomalies. *Proceedings of the 2nd ACM SIGCOMM Workshop on Internet Measurement*, 71–82. https://doi.org/10.1145/637201.637209
9. Zseby, T., Molina, M., Duffield, N., Niccolini, S., & Trammell, B. (2009). Sampling and filtering techniques for IP packet selection. *RFC 5475*. https://doi.org/10.17487/RFC5475
10. Sarrar, N., Feldmann, A., & Uhlig, S. (2012). Leveraging SPAN for network traffic analysis. *ACM Internet Measurement Conference (IMC)*, 161–174. https://doi.org/10.1145/2398776.2398794
11. Kolla, S. (2018). Legacy liberation: Transitioning to cloud databases for enhanced agility and innovation. International Journal of Computer Engineering and Technology, 9(2), 237–248. https://doi.org/10.34218/IJCET_09_02_023
12. Kim, H., & Claffy, K. (2013). Traffic characteristics of SPAN ports and TAPs in real-world deployments. *Journal of Network and Systems Management*, 21(3), 397–412. https://doi.org/10.1007/s10922-012-9248-6
13. Bolla, R., Bruschi, R., & Davoli, F. (2011). The hidden costs of port mirroring: A measurement-based analysis. *Computer Communications*, 34(9), 1080–1090. https://doi.org/10.1016/j.comcom.2010.06.005
14. Radwan, M., & Haggag, M. (2015). Comparative performance analysis of network monitoring technologies. *International Journal of Computer Science Issues*, 12(1), 56–63.
15. Varga, A., & Hornig, R. (2008). An overview of the OMNeT++ simulation environment. *Proceedings of the 1st International Conference on Simulation Tools and Techniques*, 1–10.
16. Iperf3. (2018). *Network Performance Measurement Tool*. Retrieved from https://iperf.fr/
17. Ostinato. (2018). *Packet Generator and Traffic Generator Tool*. Retrieved from https://ostinato.org/
18. Wireshark Foundation. (2018). *Wireshark Network Protocol Analyzer Documentation*. Retrieved from https://www.wireshark.org