

Rate limiting of enterprise core switch ports



Overview

CORE INTERNAL TOPOLOGY The rate limiting architecture is embedded within the Unified Forwarding Engine (UFE) ASIC.

Key Takeaways

Port rate limiting on enterprise core switches controls ingress and egress bandwidth per port, but careful planning is required to avoid unintended traffic bottlenecks, especially in multi-tenant environments.

Overview of Port Rate Limiting

Rate limiting is a mechanism to control the maximum data transfer rate on a switch port, either for incoming (ingress) or outgoing (egress) traffic. It is commonly used to:

- Prevent a single device or tenant from consuming excessive bandwidth
- Protect upstream links and network infrastructure from congestion
- Enforce service-level agreements (SLAs) in multi-tenant networks

Rate limiting can be configured per port or per VLAN, and is enforced by the switch hardware or software. Non-conforming traffic is typically dropped or shaped, depending on the switch capabilities.

Configuration Considerations



Ingress vs Egress: Ingress rate limiting controls traffic entering the switch port, while egress rate limiting controls traffic leaving the port. Both can be applied simultaneously, but the effective throughput is limited by the stricter of the two.

Hardware Enforcement: Enterprise switches often use ASIC-based enforcement for line-rate policing, ensuring that rate limits are applied without significant CPU overhead.

Multi-

Tenant Environments: When multiple tenants share a VLAN or core switch, applying different rate limits per port can lead to unexpected behavior. The lowest common denominator of configured limits may affect overall throughput if traffic is aggregated on a shared uplink. This can result in larger tenants experiencing lower-than-expected speeds if the switch cannot independently enforce per-port limits without impacting others.

Traffic Shaping vs Policing: Policing drops excess traffic immediately, which can cause packet loss for bursty applications. Shaping buffers excess traffic and smooths bursts, which can reduce packet loss but may introduce latency.

Monitoring and Verification: Use switch monitoring tools to verify that rate limits are being enforced correctly.

Compare port-level statistics with upstream carrier measurements to ensure that configured limits match actual throughput.

Best Practices

Apply rate limits at the edge of the network whenever possible, rather than on core ports, to reduce the risk of congestion affecting multiple tenants. Use consistent units (e.g., kbps or Mbps) and ensure that ingress and egress limits are aligned with the physical link capacity. For multi-tenant VLANs, consider per-tenant queuing or traffic shaping instead of strict per-port policing to allow fair bandwidth distribution.

Test configurations in a lab environment to observe the impact on bursty traffic and latency-sensitive applications. Document all rate limits and periodically review them against actual usage to adjust for growth or changes in tenant requirements.

Common Pitfalls

Aggregating multiple rate-limited ports on a single uplink can cause oversubscription, where the sum of individual limits exceeds the uplink capacity, leading to congestion and packet drops. Applying widely varying rate limits on ports within the same VLAN may result in synchronized traffic drops, as the switch enforces limits at the VLAN or hardware level. Misalignment between switch rate limits and upstream carrier bandwidth can create false bottlenecks, where the switch shows high utilization but actual internet throughput is lower.

Conclusion

Port rate limiting is a powerful tool for controlling bandwidth on enterprise core switches, but it must be applied with careful consideration of hardware capabilities, VLAN design, tenant requirements, and uplink capacity. For multi-tenant networks, per-port rate limiting should be complemented with traffic shaping, monitoring, and careful planning to avoid unintended performance degradation for larger clients.

Article Content

Enterprise Core Systems Engineering Reference Manual: Port Rate ...

CORE INTERNAL TOPOLOGY The rate limiting architecture is embedded within the Unified Forwarding Engine (UFE) ASIC. Key topological elements include: - Per-Port Policing Engine: Dedicated token

Port Rate Limiting Ingress Egress Bandwidth Deep Dive: Internal

The following table presents verified metrics from enterprise-class switches supporting port rate limiting ingress egress bandwidth with hardware acceleration and ...

Configuring Rate Limits

The traffic rate allowed by the hardware rate-limiters is configured globally and applied to each individual I/O module. The resulting allowed rate depends on the number of I/O modules in the system.

Configuring Rate Limits

Beginning with Cisco NX-OS Release 7.0(3)I6(1), you can configure a hardware rate-limiter to show statistics for outbound traffic on SPAN egress ports. This rate-limiter is supported on all Cisco Nexus

Core port rate-limiting causing issues on "tenanted" switching ...

We've ruled out local connectivity issues and want to focus here on the fact we're rate limiting to different degrees ports that are untagged in vlan 10 on the ore switch.

Enterprise Edge FAQ: Best Practices for Deploying Bandwidth Per Port ...

Q6: What is the difference between "rate-limit" and "shaping" when applying per-port bandwidth controls? Rate-limiting (policing) drops excess traffic immediately; shaping buffers excess traffic to smooth out

Port rate limiting

Configuring rate limits can help provide balance between necessary and flooded traffic. The amount of permitted traffic (the rate limit) can be specified in kilobits-per-second (kbps), packets-per-second

Port Traffic Controls

Rate-Limiting for all traffic provides a method for limiting the amount of bandwidth a user or device may utilize inbound on a switch port. This effectively sets an inbound usage level on a given port, and is a

Rate Limiting. Edge-Core ECS4620-28T, ECS4620-28F, ECS4620

The switch can set the maximum upload or download data transfer rate for any port. It can also control traffic storms by setting a maximum threshold for broadcast traffic or multicast traffic.

Example for Configuring Interface-based Rate Limiting on a Modular ...

Example for Configuring Interface-based Rate Limiting on a Modular Switch Overview
Interface-based rate limiting is easy to configure and limits the rate of all packets sent or received on

How do I configure rate limits on my NETGEAR XS512EM or XS724EM switch?

Port-based rate limiting allows you to limit the speed at which network traffic is sent or received by a device that is connected to a port on your switch. Unlike 802.1p
Quality of Service

unsupervised_topic_modeling/topics/en/15/50/100/topics at ...

Contribute to annontopicmodel/unsupervised_topic_modeling development by creating an account on GitHub.

Port based rate limiting

Hello Paul, I hate to bring up this old thread but I have a similar problem. What I am trying to accomplish is to limit the speed on some ports on Extreme Networks Switch x250 (running

QoS: Ingress Rate Limiting

Ingress rate limiting restricts the speed of traffic from a particular device to the switch port. On switches in AG mode, you must configure ingress rate limiting on F_Ports.

Switch Engine Command Reference Guide

Switch Engine™ Command Reference Guide Version 32.1 > Commands > configure ports rate-limit egress Published July 2022 Print this pageEmail this pageView PDFPreviousNext

[Enterprise Switch] How to rate limiting for Multicast DNS (MDNS)

Check the counters of multicast output packets on port 2 every second. You will find that the increase in packets is lower compared to a port without rate limiting. More information : The

Looking for a Managed Switch with easy Rate Limiting BUT...

example: on an 8-port switch, ports 2-8 have a shared rate limit of 80mbit ingress 40mbit egress, A single device on port 2 can use the total amount of that limit, Multiple devices on ports 2-8 can share

Huawei Technical Support

This guide explains how to configure interface-based rate limiting on Huawei switches, covering QoS functions like traffic shaping and congestion management.

Contact Us

Continue your research online:

Website: <https://cjcconsultants.co.za>

Technical inquiry: <https://cjcconsultants.co.za/contact.html>

This document is for preliminary research. Verify specifications against current standards, product documentation and project requirements.

