

Optimize Network Performance for Better Business Results

Improve the performance of networking, online gaming, financial trading, and intrusion-detection and prevention systems with Intel® Dynamic Load Balancer (Intel® DLB), a built-in accelerator integrated into the latest Intel® Xeon® Scalable processors.



When businesses optimize network performance and reduce latency, data moves fast. Customers can enjoy a responsive user experience, employees can be productive, and organizations can detect security threats quickly. These business benefits can lead to increased revenues and reduced costs.

To optimize network performance, data center architects and developers can turn to the latest Intel Xeon Scalable processors, which include the built-in Intel Dynamic Load Balancer (Intel DLB) accelerator. Intel offers developers streamlined tools and libraries that make it simple to unlock the potential of network infrastructure with Intel DLB. Architects trust Intel® technologies like Intel DLB to help address their top business priorities, like running apps and workloads securely, regardless of their deployment model. They can design systems with reduced latency to accelerate data across telecommunications networks. Extensive Intel partner relationships can also help architects develop sustainable solutions.

Intel DLB accelerates high-packet-rate applications by providing load balancing, prioritizing the scheduling of packets across CPU cores, and lowering latency for network traffic.

Intel DLB use cases

Intel DLB helps to distribute network traffic across multiple server CPU cores efficiently. With Intel DLB, applications like security load balancers, gateways, and content-delivery networks (CDNs)—such as streaming or forwarding—can operate with reduced latency and higher determinism for precise flow management. Functional Intel DLB use cases can include software-defined wide area networks (SD-WANs), flow policing, rate limiting, and IPsec or Transport Layer Security (TLS) gateways (see Table 1).

Table 1. Functional use cases for Intel DLB

Functional use case	Description
SD-WAN	SD-WANs dynamically route traffic over WAN connections.
Vector Packet Processing (VPP) routers	VPP routers accelerate network data-plane functions.
IPsec security gateways	IPsec gateways help establish secure communication between networks.
Elephant flow handling	Elephant flow handling helps optimize exceptionally large data flows or packets within a network.
TLS security load balancers	Security load balancers maintain proxy connections and distribute flows to free clusters.

Streamline integration

Intel makes it easy to support Intel DLB by providing the necessary libraries and software readily available in well-documented open source offerings. For example, Intel DLB communicates with data-plane software, which handles the low-level, real-time processing of data packets as they flow through the network, using standard PCI memory-mapped interfaces enabled with the [Data Plane Development Kit \(DPDK\)](#). DPDK is a collection of open source software libraries originally developed by Intel. Intel DLB can also be accessed using an Intel-provided library and kernel driver.

Intel DLB architecture

Intel DLB is a PCIe device that provides hardware solutions for managing queuing and prioritization versus common solutions that generally provide software implementations. When queue management is done in software alone (see Figure 1), many times, tolerances are wide and latencies are high. Intel DLB provides a pipelined architecture in hardware where queues are used to connect various software stages. Load balancing and complex scheduling might be required between stages to drive improved network performance.

Queue management: before and after

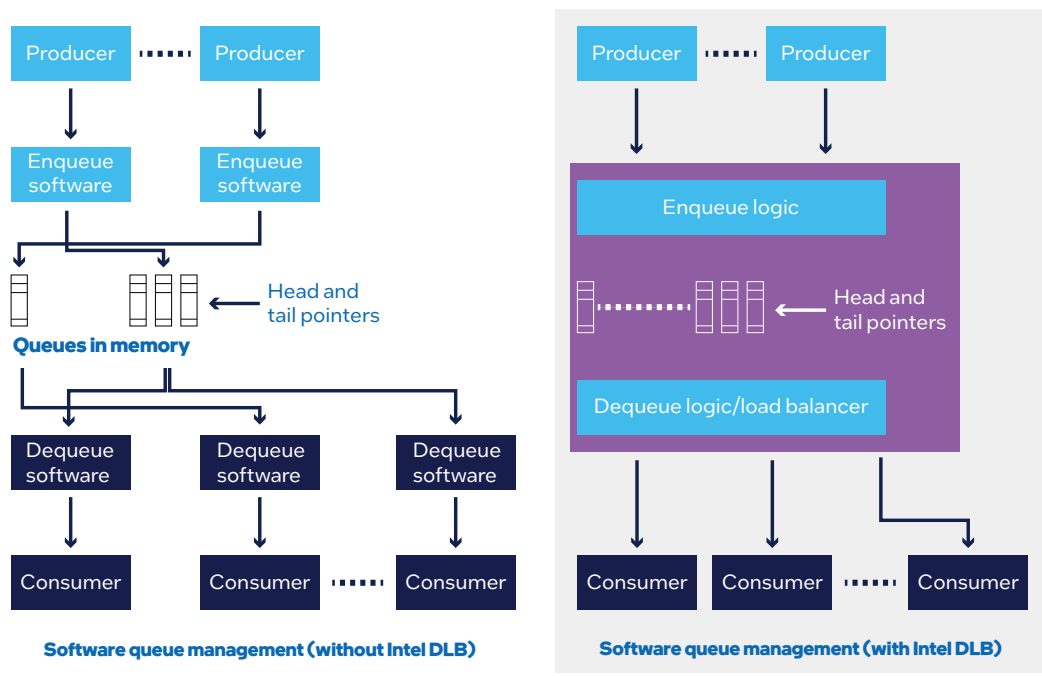


Figure 1. Intel DLB supports queue management in hardware

Intel DLB’s hardware implementation provides several key benefits:

- Queue management is offloaded from software to hardware.
- Load balancing is dynamic and flow-aware.
- It supports up to eight levels of priority queuing.
- It allows dynamic sizing of applications for seamless scaling up or down.

Developer tools

Intel offers developers and solution architects a variety of tools to take advantage of Intel DLB:

- [Intel DLB software user guide](#): Provides configuration and enablement details for Intel DLB
- [Intel® Developer Zone](#): Official source for developing software on Intel hardware

- [Intel® Developer Cloud](#): Resources to help developers learn, prototype, test, and run workloads on a cluster of the latest Intel hardware
- [Intel® DevCloud for oneAPI](#): Development sandbox for learning about programming cross-architecture applications
- [Driver for Intel DLB](#): DPDK documentation for the poll-mode driver used to configure devices and their respective queues; the driver works with Intel DLB software versions 2.0 and 2.5
- [Intel DLB Linux release download](#): Downloadable version of Intel DLB for Linux
- [Proof points of Intel DLB](#): Technical information about Intel DLB

Enable Intel DLB for key workloads

Developers and architects can get started with Intel DLB by using the following resources and tools, which can make it easy to implement Intel DLB in existing solutions:

- [Intel DLB scaling of IPsec workloads technology guide](#): Learn how to use Intel DLB as a hardware queue manager and load balancer to offload queuing and scheduling tasks from software for IPsec workloads.
- [Intel DLB telemetry/insight demo](#): Explore how to enable and use Intel DLB for telemetry workloads.

Accelerate network performance

The following scenarios illustrate how the latest Intel Xeon Scalable processors with Intel DLB can accelerate Istio, Envoy, NGINX, microservices on Kubernetes, and IPsec elephant flow applications. For further reading, workload-specific resources are provided.

Istio and Envoy

Istio and Envoy are two open source projects that are often used together to provide advanced networking and service-mesh capabilities in modern containerized and microservices-based applications.

Use 4th Gen Intel Xeon Scalable processors with Intel DLB to achieve up to 96 percent lower latency at the same throughput (RPS) with 4th Gen Intel Xeon Scalable processors and Intel DLB versus software for an Istio ingress gateway working on 6 cores/12 threads.¹

Enablement resources:

- [Accelerated Offload Connection Load Balancing in Envoy; Envoy's First Hardware Feature](#): Information about how Intel DLB accelerates connection load balancing in Envoy
- [Using Accelerated Offload Connection Load Balance in Istio](#): Documentation for using Intel DLB connection load balancing in Istio
- [Istio and Envoy optimizations for Intel Xeon Scalable processors](#): Solution brief for Istio and Envoy integration

Container details:

- [Envoy Docker Hub](#): Information about the Envoy proxy for Docker containers

Coding resource:

- [Envoy DLB connection balancer](#): Documentation (including code) for integrating Intel DLB into Envoy

NGINX

Intel DLB benefits NGINX web server software, which is used by 34 percent of websites.² Software architects and developers can enhance the NGINX HTTP2/3 distributor using Intel DLB to offload the distribution of requests across worker cores at the server. This is especially helpful when different HTTP requests have different priorities. A hardware-based Intel DLB priority queue can also be used in an NGINX thread pool.

Use 4th Gen Intel Xeon Scalable processors with Intel DLB to provide latency improvements of up to 1.3–2.3x when queuing is offloaded to Intel DLB for various object sizes.³

Enablement resource:

- [Microservices Solution – Optimizations with Intel® Xeon® Scalable Processor](#): Solution brief about performance optimizations for data center operators and microservices developers working with NGINX

Microservices on Kubernetes

Microservices on Kubernetes refers to a software architecture and deployment approach where microservices-based applications are containerized and managed within a Kubernetes cluster. Software architects and developers can get results faster by using caching for microservices with a Memcached workload.

Use 4th Gen Intel Xeon Scalable processors with Intel DLB to reduce latency up to 22–42 percent and cycle utilization up to 30–60 percent on a server.³

Enablement resource:

- [Microservices Solution – Optimizations with Intel® Xeon® Scalable Processor](#): Solution brief about performance optimizations for data center operators and microservices developers

Container details:

- [Intel DLB device plugin for Kubernetes](#): A framework for developing Intel DLB plugins for Kubernetes

IPsec elephant flow

Software-based network gateways face the challenge of elephant flow, the transmission of exceptionally large amounts of data over a network. When the volume of the flow exceeds the processing capability of the system, packets might be dropped, which can impact the quality of service (QoS).

Enablement resource:

- [Intel® Dynamic Load Balancer \(Intel® DLB\) - Accelerating Elephant Flow](#): Technology guide that provides documentation about the coordination and collaboration of multiple CPU cores to handle elephant flows

Intel DLB with 4th Gen Intel Xeon Scalable processors used for elephant flow handling and linear scalability achieves a maximum of up to 22 Mpps processing throughput for a single flow.⁴

Optimize network performance

Intel DLB has the potential for very-high-rate packet scheduling and queuing to optimize network performance and move data faster. This can lead to more efficient business processes, faster responses to security threats, and—as a result—happier customers.

Learn more about [Intel Xeon Scalable processors](#) and [Intel® Accelerator Engines](#), and start enabling Intel DLB today.



¹ See claim [W6] at www.intel.com/PerformanceIndex; 4th Gen Intel Xeon Scalable processors. Results may vary.

² W3Techs. "Usage statistics of Nginx." September 2023. <https://w3techs.com/technologies/details/ws-nginx>.

³ Intel. "Microservices Solution – Optimizations with Intel® Xeon® Scalable Processor." February 2023.

<https://networkbuilders.intel.com/solutionslibrary/microservices-solution-optimizations-with-intel-xeon-scalable-processor-solution-brief>.

⁴ Intel. "Intel® Dynamic Load Balancer (Intel® DLB) - Accelerating Elephant Flow." February 2023.

<https://networkbuilders.intel.com/docs/networkbuilders/intel-dynamic-load-balancer-intel-dlb-accelerating-elephant-flow-technology-guide-1677672283.pdf>.

Performance varies by use, configuration and other factors. Learn more at www.intel.com/PerformanceIndex.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See configuration disclosure for additional details.

No product or component can be absolutely secure.

Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.