

CONSISTENTLY MEET SERVICE LEVEL AGREEMENTS

Meet the challenges of predictability at scale – and consistently meet SLA commitments – with Intel® Ethernet 800 Series with Application Device Queues (ADQ)

Throughput, Latency – and Predictability

End users expect network responses within a given amount of time. As data centers scale, application response times become more variable, even in fast and efficient systems. You may not be able to consistently meet service level agreements (SLAs) and you risk disappointing users. The network plays a key role in influencing the predictability of application response time.

Express Lanes for Data

Like cars on a highway, data needs to travel from origination to destination quickly and efficiently, but often gets stuck with too much traffic and unexpected delays. Intel® Ethernet 800 Series has ADQ that acts as dedicated express lanes for high-priority applications.

General purpose traffic can be unpredictable.

GENERAL PURPOSE LANES

2,048 ADQ EXPRESS LANES

ADQ filters application traffic to a dedicated set of queues for more predictable high performance.

<EROSPIKE


redis

MEMCACHED

SLA SUCCESS

Proven Performance

ADQ accelerates Ethernet network performance, providing increased application response time predictability, lower latency, and higher throughput.



Reliably Meet SLAs

Intel Ethernet 800 Series with ADQ helps you more consistently meet SLAs. High predictability means you can use more servers to parallelize a task and process requests more quickly. Or, with ADQ, your existing servers can support more end users and can process more requests for lower TCO.

Intel is dedicated to database optimization, helping organizations move, store, and process data. Find out how Intel Ethernet 800 Series with ADQ can help your business optimize server performance, enabling greater predictability and helping you confidently meet your SLAs.

LEARN MORE ABOUT INTEL® ETHERNET 800 SERIES WITH ADQ

[VISIT INTEL.COM/ETHERNET](https://www.intel.com/ethernet)

Performance results are based on Intel internal testing as of February 2019, and may not reflect all publicly available security updates. See configuration disclosure for details. No product or component can be absolutely secure. For complete configuration information see the [Performance Testing ADQ with Redis Solution Brief](#).

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software, or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer, or learn more at [intel.com](#).

Performance results are based on testing as of the date noted in the configuration details and may not reflect all publicly available security updates. See configuration disclosure for details. No product or component can be absolutely secure. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit [intel.com/benchmarks](#).

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit [http://www.intel.com/performance](#). Source: Performance results are based on Intel internal testing as of February 2019, and may not reflect all publicly available security updates. See configuration disclosure for details. No product or component can be absolutely secure. Tests performed using Redis Open Source on 2nd Generation Intel® Xeon® Scalable processors and Intel® Ethernet 800 series 100GbE on Linux 4.19.18 kernel. Calculation: (new - old) / old x 100%. For reduction in Variance of Standard Deviation of Rtt Average Latency across all runs (10 to 100) for baseline vs ADQ (229-739)/739 * 100% = -69% Reduction in Variance; Rtt Average Latency across all runs for baseline vs ADQ (382-1249)/1249 * 100% = -69% Reduction in Rtt Average Latency; Transaction Request Rate across all runs for baseline vs ADQ (79601-44345)/44345 * 100% = 80% Throughput Improvement.

Intel, the Intel logo, and Xeon are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries. © Intel Corporation

* Other names and brands may be claimed as the property of others. 0520/TKOE/KC/PDF

