

Index up to 2.49 Times as Many Documents per Second and Shorten Search Times for Elasticsearch Workloads with Microsoft Azure Ddsv5 VMs Featuring 3rd Gen Intel® Xeon® Scalable Processors



Elasticsearch



Index up to 2.49 times as many documents per second on Microsoft Azure Ddsv5 VMs with 3rd Gen Intel Xeon Scalable processors

vs. Dsv3 instances



Reduce search times by 69% on Microsoft Azure Ddsv5 VMs with 3rd Gen Intel Xeon Scalable processors

vs. Dsv3 instances


Ddsv5 VMs Sped Up Indexing Over Both Dsv3 and Ddsv4 VMs with Previous-Generation Processors

Microsoft Azure Kubernetes Services (AKS) offers flexibility and agility to organizations running Elasticsearch workloads on Kubernetes clusters in the cloud. Which VM type offers the best level of performance for these workloads? To find out, we compared the Elasticsearch performance of three instance types in a Kubernetes containerized environment:

- Microsoft Azure Ddsv5 VMs with 3rd Gen Intel® Xeon® Scalable processors
- Microsoft Azure Ddsv4 VMs with 2nd Gen Intel Xeon Scalable processors
- Microsoft Azure Dsv3 VMs with Intel Xeon E5-2673 v4 processors

To quantify Elasticsearch performance, we used the Rally 2.3.0 benchmark tool, which acts as a load generator and records and stores telemetry metrics in Elasticsearch during search workloads to find performance problems for comparison across clusters. Using two datasets from Rally, `nyc_taxi` (a transport analytics dataset) and `http_logs` (a Web services dataset), we found that Azure Ddsv5 VMs with 3rd Gen Intel Xeon Scalable processors indexed up to 2.49 times as many documents per second with as much as 69 percent shorter search times vs. an Azure Dsv3 VM cluster. Additionally, the Ddsv5 VMs outperformed the Ddsv4 VMs by up to 30 percent in indexing throughput with 29 percent shorter search times. For more information about Rally and the workloads we used, see below.

This shows that across multiple Elasticsearch workloads, selecting Azure Ddsv5 VMs with the latest processor technology can offer better Azure Kubernetes Services performance than VMs built on two previous processor generations. By indexing content faster, organizations get the data they need in less time to make vital business decisions.



Index Data Faster to Gain Quicker Insights with Azure Ddsv5 VMs

While organizations use Elasticsearch for many use cases, including application performance monitoring, application searches, and business analytics, we compared the indexing throughput for a transport analytics workload and a typical Web services workload (see Figure 1).

Using the nyc_taxi dataset, which stores 2015 NYC taxi ride data and runs search scenarios, Ddsv4 VMs with 2nd Gen Intel Xeon Scalable processors indexed 53 percent more documents/second (docs/s) than the Dsv3 VM cluster. Azure Ddsv5 VMs with 3rd Gen Intel Xeon Scalable processors improved performance even more, offering 91 percent more docs/s than the Dsv3 VMs.

On the http_logs dataset, which uses example HTTP server log data and runs search scenarios, the VM types again showed similar performance scaling over the Dsv3 VMs. Azure Ddsv4 VMs had 92 percent more throughput, while Ddsv5 VMs enabled by 3rd Gen Intel Xeon Scalable processors had 2.49 times as much throughput, for significantly faster indexing.

Normalized Indexing Throughput

Documents per second | Higher is better




Figure 1. The relative indexing throughput, in documents per second, that the clusters achieved running two types of datasets from the Rally benchmark to assess Elasticsearch performance. Higher numbers are better.

Keep Search Times Down with Azure Ddsv5 VMs

The Rally benchmark also captured search latency while running Elasticsearch workloads. Shorter search times mean that organizations don't have to wait as long to gain insights from their data. Figure 2 compares the normalized search latency the VMs experienced for both the nyc_taxi and http_logs datasets. Ddsv4 VMs with 2nd Gen Intel Xeon Scalable processors reduced search latency by 33 percent and 44

percent, respectively. The latest Ddsv5 VMs with 3rd Gen Intel Xeon Scalable processors shrunk wait times even further, reducing search latency by 56 percent and 69 percent, respectively, compared to the Dsv3 VMs.

Normalized Search Latency

Milliseconds | Lower is better




Figure 2. The relative search latency, in milliseconds, that the clusters achieved running two types of datasets from the Rally benchmark to assess Elasticsearch performance. Lower numbers are better.

Conclusion

No matter your particular use case for Elasticsearch, indexing more content per second can lead to quicker insights. Tests show that selecting Microsoft Azure Ddsv5 VMs featuring 3rd Gen Intel Xeon Scalable processors to run Elasticsearch on Kubernetes clusters can improve indexing throughput and search times for multiple use cases.

If the cloud instance type you choose can index more data in less time for faster data insights, it may also reduce the number of instances you need to purchase, potentially improving your operating costs.

Learn More

To begin running your Elasticsearch workloads on Microsoft Azure Ddsv5 VMs, visit <https://docs.microsoft.com/en-us/azure/virtual-machines/ddv5-ddsv5-series>.

To learn more about Elastic, visit <https://www.elastic.co/>.

To use the Rally benchmark tool, visit <https://github.com/elastic/rally>. You can find both datasets used in testing at https://github.com/elastic/rally-tracks/tree/master/http_logs and https://github.com/elastic/rally-tracks/tree/master/nyc_taxis.

Tests by Intel as of 12/29/21. All configs consisted of 3 VMs with 16 cores, HT On, Turbo Off, Total Memory 64 GB, Hyper-V UEFI Release v4.1, 0xfffffff, Ubuntu 18.04.6 LTS, 5.4.0-1064-azure, N/A, elastic/elasticsearch:7.16.2, elastic/rally:2.3.0, run_type: nyc_taxi/http_logs, java_opts: -Xms12g -Xmx12g number_of_shards: 36 number_of_replicas: 1. D16s_v3: Intel(R) Xeon(R) CPU E5-2673 v4 @ 2.30GHz, D16ds_v4: Intel(R) Xeon(R) Platinum 8272CL CPU @ 2.60GHz, D16ds_v5: Intel(R) Xeon(R) Platinum 8370C CPU @ 2.80GHz



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 backup for configuration 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 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.

Printed in USA 0422/JO/PT/PDF US001

