



Bringing Scalability to Metadata Searching

phoenixNAP worked with Vizion.AI to launch phoenixNAP Elasticsearch* Service, which enables huge volumes of metadata to be searched and stored in object storage. Intel® Optane™ DC persistent memory accelerates the cache

At a glance:

- Companies are struggling to manage the growing volumes of machine-generated data (metadata).
- Elasticsearch* presents challenges in managing server clusters and usually requires data to be stored in flash storage, which becomes expensive as data volumes increase.
- Using Vizion.AI's technology, it's possible to store the data in highly scalable object storage and use Intel® Optane™ DC persistent memory as a cache.
- By launching the phoenixNAP Elasticsearch Service, phoenixNAP has created new opportunities for differentiation and growth.
- Intel worked with phoenixNAP and Vizion.AI to accelerate the launch of the service.

Data centers and the rising numbers of connected devices generate a huge amount of metadata. Companies need an effective way to store and search it, to help them manage their infrastructure and detect security breaches. Using phoenixNAP Elasticsearch* Service, companies can now store their data in scale-out object storage instead of memory and keep a cache in Intel® Optane™ DC persistent memory to accelerate performance.

Challenge

- Storing and analyzing the growing volumes of metadata can be challenging, yet it is necessary for managing the IT estate and its security.
- Logs may be distributed across different cloud platforms, making it difficult to build a holistic picture of a particular event or user journey.
- Elasticsearch provides a solution, but the clustering requires ongoing management and data is typically kept in flash storage, which may be prohibitively expensive as data volumes rise.

Solution

- phoenixNAP Elasticsearch Service enables companies to easily deploy Elasticsearch to search across data in different cloud locations.
- Using Vizion.AI's technology, the service enables affordable object storage to be used for 90 percent of the data.
- The other 10 percent is a cache of the hottest data, which is stored in Intel Optane DC persistent memory, providing near-memory speeds at affordable large capacities.

Results

- Using Intel Optane DC persistent memory for the cache cut latency by 80 percent and accelerated indexing by 3x compared to hosting the solution in a hyperscale cloud environment¹.
- phoenixNAP customers can now analyze the metadata stored in phoenixNAP's data centers without needing to validate the data sovereignty of a new provider.
- The cloud service provider (CSP) can also sell the service to new customers that need an easier way to make sense of their metadata.



Managing huge volumes of metadata

By 2025, the Internet of Things (IoT) will be generating 79.4 zettabytes of data a year, according to IDC². That's roughly 79.4 billion terabytes. Along with all the data people want comes metadata: data about the data, such as logs. That will play a role in this data growth and will need to be managed and analyzed. This will intensify a problem that companies already face: many digital assets, including servers and applications, are generating logs that they must analyze to manage their IT performance and help to protect its security.

These logs might be distributed across different cloud platforms, and finding any single data point can mean huge volumes of data need to be searched. Often multiple data points will need to be correlated. If there is a security incident, for example, companies might need to quickly trace the activities of a user across all the logs within a particular time period. The longer that takes, the greater the risk to the business.

Many companies are looking for a cloud-native platform that will enable them to not only store, but also to search and analyze, their growing metadata logs. Elasticsearch is an open source search engine that could be used, but it typically requires the data to be stored on flash storage media, which may be unsustainably expensive as data volumes grow. Additionally, deployment and management of Elasticsearch can be demanding. Customers might need to launch six virtual machines before they can get their first data in and will need to monitor data volumes so they can delete data or expand the cluster when it fills.

Putting the data into a hyperscale cloud environment may be cheaper, but it still requires ongoing cluster management by the user and does not take advantage of low-cost object storage.

Cloud customers are looking for a solution that overcomes these limitations, to easily and quickly analyze their metadata, while taking advantage of the economics of open source software.

Introducing phoenixNAP Elasticsearch* Service

CSP phoenixNAP worked with Vizion.AI and Intel to launch the phoenixNAP Elasticsearch Service. This service enables companies to easily deploy Elasticsearch to analyze data across their multi-cloud environment. Only the hottest 10 percent of data needs to be stored in fast storage: the rest can be kept in the cloud using object storage, with Vizion.AI's solution taking care of compression, deduplication, encryption, and transport to and from the cloud. Using the cloud for the bulk of the metadata can dramatically cut the cost, compared to keeping it in flash storage.

Customers order the solution through a portal and the set-up and management is automated, resolving one of the pain points with Elasticsearch. Because the solution is based on

a microservice rather than a managed infrastructure pool, customers no longer need to monitor their data volumes. The object storage on the back-end scales to accommodate data as it comes in.

Figure 1 shows a simplified architecture of the solution. Customer workloads run in containers and are orchestrated using Kubernetes*. Because the infrastructure is shared across hundreds of clients, there are significant economies of scale, compared to a company setting up its own Elasticsearch cluster.

Vizion.AI's parent company Panzura* provides an intermediary layer that can translate between data center storage protocols, and cloud-native object storage. This enables Elasticsearch (and other applications) running in the data center to access cloud storage without modification, with the in-memory cache helping to deliver high performance.

The underlying server hardware is powered by the 2nd generation Intel® Xeon® Scalable processor family with Intel Optane DC persistent memory. Intel Optane DC persistent memory enables CSPs to unlock a unique combination of affordable large capacities, with near memory performance. Physically, Intel Optane DC persistent memory is compatible with DRAM and plugs into the same DIMM slots. The persistent memory is used for caching the hottest 10 percent of data.

The virtualization layer is based on VMware vSphere* 6.7, which supports Intel Optane DC persistent memory.

Vizion.AI measured the performance of the infrastructure using persistent memory and compared it to the speed of its software running in a public cloud service. The company found that the phoenixNAP implementation was 3x faster at indexing and digitizing documents³. There was also an 80 percent reduction in latency, which is particularly important when using any identified issues to trigger an incident response in real time⁴.

"The nice thing about using Intel Optane DC persistent memory is that it's not disruptive to app deployment," said Geoff Tudor, Vice President and General Manager of Vizion.AI. "We look at the world from a Kubernetes perspective, and persistent memory looks to us just like another storage resource. We get efficiencies out of the box, without having to change any code. Having our solution accelerated by Intel Optane DC persistent memory enables us to achieve cost efficiencies too, compared to using a hyperscale cloud provider."

With several hundred customer containers on each server, improvements in the indexing of new content and latency in real-time search help avoid contention on the server and improve the customer experience.

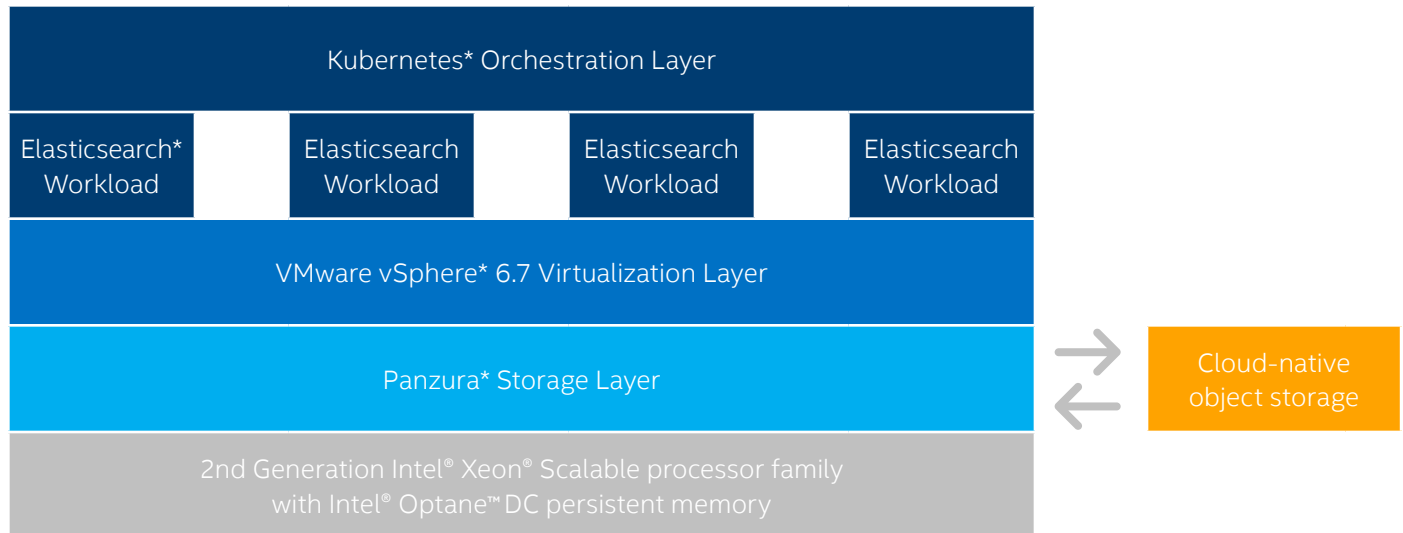


Figure 1. The phoenixNAP Elasticsearch* Service enables object storage to be used for scale-out storage of metadata which can be searched and analyzed using Elasticsearch. The solution is based on the 2nd generation Intel® Xeon® Scalable processor family with Intel® Optane™ DC persistent memory.

Technical Components of the Solution

- **Elasticsearch*.** Elasticsearch is open source software that provides a distributed, multitenant-capable full text search engine.
- **Panzura* Storage Layer.** The Panzura Storage Layer provides an intermediary between data center storage protocols and cloud storage protocols, enabling data center applications to use cloud storage without modification.
- **2nd generation Intel® Xeon® Scalable processor.** The 2nd gen Intel Xeon Scalable processor provides the foundation for a powerful data center platform that creates a leap in agility and scalability. Disruptive by design, this innovative processor sets a new level of platform convergence and capabilities across compute, storage, memory, network, and security. Enterprises and cloud and communications service providers can now drive forward their most ambitious digital initiatives with a feature-rich, highly versatile platform.
- **Intel® Optane™ DC persistent memory.** Intel Optane DC persistent memory is a new class of memory that brings greater capacity to the cores – terabytes instead of gigabytes per platform – and is accessible over the memory bus. This revolutionary technology delivers a unique combination of affordable large capacity and support for data persistence. It is supported by 2nd gen Intel Xeon Scalable processors.

Intel is an Ally

Intel has a close working relationship with phoenixNAP, including helping the company to develop and market new services. To accelerate the launch of phoenixNAP Elasticsearch Service, Intel provided phoenixNAP with early access to Intel Optane DC persistent memory, and the 2nd gen Intel Xeon Scalable processor which is required to use it. Intel was on hand to offer support with implementing the solution, and with fine-tuning it to improve performance.

Differentiating with Elasticsearch*

The cloud market is intensely competitive, so phoenixNAP differentiates from hyperscale providers by offering a cloud environment that is optimized for particular applications and enhanced with value-added services.

The launch of the new Elasticsearch service enables phoenixNAP to provide more value to its existing customers, by giving them new ways to manage their data within the phoenixNAP cloud. For customers subject to regulation, that means there's no need to validate the data sovereignty or security of a new solution provider.

The new service can also help to attract new business. Because the solution works across cloud environments, it does not require the bulk of the data to be stored in phoenixNAP's data center. New customers may choose to migrate data to phoenixNAP, or may prefer to take advantage of the ease of deployment phoenixNAP offers as they access third-party cloud storage locations.

Vizion.AI and phoenixNAP are working together on marketing the solution to build their joint customer base.

Lessons Learned

There are several lessons from phoenixNAP and Vizion.AI's experience launching the new service.

- “Matching the workload to an optimized infrastructure can deliver huge savings,” said Geoff Tudor. “Using Intel® Optane™ DC persistent memory, I can get more processing done than I could on an equivalent hardware platform without persistent memory.”
- Code modifications are not always required to benefit from Intel Optane DC persistent memory. VMware vSphere* 6.7 is compatible with persistent memory, for example, out of the box.
- Intel works with cloud service providers (CSPs) to help them to create and launch new services, including the new phoenixNAP Elasticsearch* Service.

Find the solution that is right for your organization. Contact your Intel representative or visit intel.com/csp


Spotlight on phoenixNAP

Founded in 2009, phoenixNAP is a global IT services provider offering cloud, dedicated server, colocation and infrastructure as a service (IaaS) technology solutions. phoenixNAP is a Premier Service Provider in the VMware vCloud Air* Network Program and is a Payment Card Industry Data Security Standard (PCI DSS) Validated Service Provider. Its flagship facility in Phoenix, Arizona, is Service Organization Controls (SOC) Type 1 and SOC Type 2 audited.

For more information, visit phoenixnap.com

Learn More

- 2nd generation Intel® Xeon® Scalable processors
- Intel® Optane™ DC persistent memory



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 product or component can be absolutely secure. Check with your system manufacturer or retailer or learn more at intel.com.

^{1,3,4} Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors.

² <https://www.idc.com/getdoc.jsp?containerId=prUS45213219>

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 www.intel.com/benchmarks

Configurations: Up to 3x indexing and 80% cache latency decrease – based on phoenixNAP and Panzura testing as of March 2019 on Elasticsearch: Intel® Xeon® Gold 6230 processor, Total Memory 256 GB RAM, 1.5TB of Intel® Optane™ DC persistent memory, HyperThreading: Enabled, Turbo: Enabled, ucode: 0x043, OS: [‘centos-release-7-5.1804.el7.centos.x86_64’], Kernel: [3.10.0-862] vs. AWS i3xlarge (Intel) Instance, Elasticsearch, Memory: 30.5GB, Hypervisor: KVM, Storage Type: EBS Optimized, Disk Volume: 160GB, Total Storage: 960GB, Elasticsearch version: 6.3

Performance results are based on testing as of the date set forth in the Configurations and may not reflect all publicly available security updates. See configuration disclosure for details. No product or component can be absolutely secure.

Intel does not control or audit third-party data. You should review this content, consult other sources, and confirm whether referenced data are accurate.

Cost reduction scenarios described are intended as examples of how a given Intel-based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

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

*Other names and brands may be claimed as the property of others.