



Optimizing IT Operational Analytics to Empower a Digital Business—The Mandate to Go Beyond “Good Enough”

IT teams need to analyze ever-increasing amounts of log data, security data and application performance data to ensure digital systems are really protecting and supporting the business. Most organizations are currently running analytics systems using Amazon OpenSearch Service to support these use cases. However, many have not yet optimized these workloads' functionality and performance. Many organizations have relied on initial implementations because OpenSearch just keeps working. Inefficient implementations may be good enough today, but optimizing them can both reduce costs and free up resources that can be used for other projects. This eBook uses external deployment data from AWS to identify simple steps any IT team can take to dramatically enhance the efficiency of their OpenSearch Service.

Chapter 1
Introduction 3

Chapter 2
Overview of Amazon OpenSearch Service 5

Chapter 3
More Bang for the Buck—Amazon OpenSearch Service
Delivers Enhanced Cost Efficiency 6

Chapter 4
Securing Critical Data in OpenSearch Service 8

Chapter 5
Enhancing Availability and Redundancy 10

Chapter 6
Tying It All Together for Better Outcomes 12

Chapter 7
Conclusion and Next Steps 13

Chapter 1

Introduction

IT teams face a clear mandate—keep the organization’s digital systems up, responsive and secure. IT teams are utilizing critical log and event data from Amazon OpenSearch Service to gain visibility into these workloads to ensure these challenges are met. It’s how IT teams keep users happy and operations humming. This service makes leading-edge log analytics, application monitoring, security information and event management (SIEM) and document search possible.

Organizations benefit in several ways by optimizing their OpenSearch Service implementation. While current implementation might be operating at an acceptable level to meet current needs, new workloads and use cases are always being added. This may result in inefficiencies or sub-optimal results in the future. It is imperative to stay current and leverage new features and capabilities to improve cost efficiency, support current demands and enable IT teams to solve future problems.

An analysis of current OpenSearch Service usage patterns provides insight into specific changes that should be strongly considered. Amazon OpenSearch Service currently has tens of thousands of active customers, with hundreds of thousands of clusters under management, processing trillions of requests per month. Anonymized service data is used by AWS to identify the most beneficial changes to make. The chart on the following page showcases some of the key actions IT teams can take to meet business demands today and tomorrow.

IT teams face a clear mandate—keep the organization’s digital systems up, responsive, and secure. IT teams are utilizing critical log and event data from Amazon OpenSearch Service to gain visibility into these workloads to ensure these challenges are met. While open source solutions like OpenSearch offer a powerful set of features that simplify deploying an enterprise-grade application, IT teams need to ensure they are managing security, ensuring scalability, monitoring and tuning the deployment to get the best results.



Optimizing Compute & Storage	Optimizing Security	Optimizing Availability/ Resiliency
<ul style="list-style-type: none"> • Move to Graviton2 instances that offer 30% price/performance improvement. • Leverage the efficiency of Graviton instances that are 10% lower in cost than 5th-generation instances (only 15% of users are running Graviton). • Move to newest generation of instances. This can reduce costs by 68% (88% of customers aren't fully leveraging them). • Utilize GP3 Storage, as it offers 10% cost reduction and 1.5x density increase (less than 5% of customers use GP3). • Deploy cold and UltraWarm storage, as these options can reduce costs 80% (less than 10% of customers use it). 	<ul style="list-style-type: none"> • Implement node-to-node encryption to improve protection (only 55% of customers use it). • Utilize fine grained access control, as it offers broad protection (it is used by only ~25% of customers). • Move to the latest versions of OpenSearch, as they offer many security enhancements (80% of customers can improve security posture simply by moving to the latest version). 	<ul style="list-style-type: none"> • Follow AWS's best practices and guidelines to get 99.9% availability. • Utilize three Availability Zones (AZs) to dramatically improve resilience (less than 20% of customers are using three). • Utilize the latest versions with the most up-to-date software, complete with bug fixes and continual stability optimization, to improve availability (less than 43% of customers use these versions).

Getting the Most Out of Amazon OpenSearch Service: Cost Efficiency, Security, Resiliency

OpenSearch Service users will often want to optimize their use of OpenSearch analytics and management tools to reduce costs, improve security and increase reliability. The data provided in this eBook will provide readers with an understanding of their options, the benefits of different approaches and other key details to inform key decisions for this critically important technology solution. The research data makes it possible to more easily compare the impact of the different options that are available.

One important consideration that needs to be part of the evaluation process is the use of managed services to offload operational tasks from the existing IT team and to simplify the migration to the latest generation of tools and solutions for these critical activities. Historically, upgrades and enhancements were almost always self-managed, but as the strategic demands on IT move toward a strategic partnership, leading firms are finding that the use of managed services will free up IT resources to better align with that goal.

A key feature of the AWS solution is the delivery of non-disruptive upgrades and new versions of the service. The problem of upgrades that break current systems and processes has been a key reason why organizations don't move to the latest open source offerings. However, even with seamless upgrades across minor versions and clear guidance to move across major versions, customers are slow to change. While customers readily use the latest versions for their new domains (demonstrated by majority of new domains being created with the latest versions in OpenSearch Service), the adoption is slower when it comes to customers moving their existing domains to more recent versions. While moving across major versions might involve some incremental effort, the cumulative value that customers can derive from new capabilities will outweigh the incremental effort required in moving versions. It is simpler when customers do not skip multiple major versions but rather invest in upgrading their domains on a regular basis. When customers have to upgrade from a much older version (e.g., three major versions apart) to a new version, the effort will be higher and the testing more complex due to the changes across multiple versions. ■

Chapter 2

Overview of Amazon OpenSearch Service

Amazon OpenSearch Service is designed to perform real-time search, monitoring and analysis of operational and log data. OpenSearch Service is an AWS managed service that allows IT teams to run open source OpenSearch clusters without investing in managing, monitoring or maintaining the software or infrastructure to support it. This service supports both OpenSearch and Elasticsearch engines.



To simplify deployment and use, AWS manages the software installation, upgrades, patching, scaling and cross-Region replication. This service is also bundled with a dashboard tool, OpenSearch Dashboards, to make it easy to visualize results for log and trace data, in addition to supporting machine learning-powered results for anomaly detection.

AWS is regularly adding features and functionality to meet the evolving

customer demands on this solution. Some examples of recent AWS-delivered enhancements include cross-cluster replication, trace analytics, transforms and notebooks for OpenSearch Dashboards. An important advantage of the AWS offering is the focus on reducing problems or breakage that occur when new versions of the solution are deployed. Seamless upgrades are now possible to provide additional and new functionality that is simpler and more expedient.

Perhaps the most attractive aspect of utilizing OpenSearch Service is AWS's focus on enhancing the open source engines with the complimentary technologies and solutions that deliver more capability and business value for IT teams. Self-managing open source solutions refer to exactly that: You manage the security and administration of your installation. A fully managed service removes the management, administration, integration, and infrastructure security demands, reducing the demand on an organization's IT team. AWS has focused on delivering a complete service offering to ensure efficiency, reliability, and security. In addition, OpenSearch Service includes integrations with several other AWS services automatically, like AWS CloudTrail, AWS CloudWatch, AWS EventBridge, and more. ■

Some examples of recent AWS-delivered enhancements include cross-cluster replication, trace analytics, transforms and notebooks for OpenSearch Dashboards. An important advantage of the AWS offering is the focus on reducing problems or breakage that occur when new versions of the solution are deployed.


Chapter 3

More Bang for the Buck—Amazon OpenSearch Service Delivers Enhanced Cost Efficiency

In the past, log analytics, application monitoring and security analytics were done on only a subset or small number of workloads because it takes time and energy to implement them successfully. However, the current IT environment requires more comprehensive analytics and monitoring across most—and, in some cases, all—workloads. This has resulted in increased costs, and the inefficiencies of older technology and legacy processes for these activities are no longer affordable.

Improving Compute Cost Efficiency

There are several ways that organizations can improve the efficiency of the compute instances and resources that are being used. AWS is consistently enhancing price performance of new compute instances. A good example of this is Graviton. Each new generation of an instance family has not only raised the bar in terms of performance, but has also reduced the cost of running the service. In some instance families, the price has been reduced by as much as 68% from generation to generation, while the performance has improved by up to 70%. For example, the latest generation of Graviton instances offer up to a 30% price-to-performance improvement over the 5th-generation x86 instances. Instances are available for as low as \$0.036/hr. However, according to AWS's analysis of current usage patterns, 88% of OpenSearch Service customers have yet to take advantage of the cost efficiency possible from simply moving to the most current instances. In addition, OpenSearch Service offers Reserved Instances, which provide up to 52% savings annually off the list price. Surprisingly, AWS's data shows that this cost reduction option is often overlooked, as only 18% of customers leverage this pricing option.



Among AWS OpenSearch Service customers, Graviton instances are the fastest-growing instance type for the managed OpenSearch Service. Some key characteristics of the most up-to-date (m6g.large) Graviton instances include the following:

- They are 10% less expensive (\$0.128) than 5th-generation X86 (m5.large - \$0.142).
- They are 17% less expensive than 4th-generation X86.
- They are 32% less expensive than 3rd-generation X86.

Yet, the AWS analysis shows that only 15% of customers are currently using Graviton instances.

...according to AWS's analysis of current usage patterns, 88% of OpenSearch Service customers have yet to take advantage of the cost efficiency possible from simply moving to the most current instances.

Another important way to optimize this service is to choose the most efficient compute instance from a range of different instance types. Instances are focused on specific types of workloads, including, general-purpose, compute-optimized, memory-optimized or storage-optimized. Considering that the OpenSearch Service can support several use cases, running this managed service on the most effective instance for a specific use case or workload improves efficiency.

Improving Storage Efficiency

The cost of cloud storage is an important element of the overall cost equation for cloud services, and, more importantly, optimizing storage can delivery cost savings. Storage costs for databases or storage-heavy applications in OpenSearch should be managed closely. To start, there are many new and different options for cloud storage that are now available. In much the same way that matching compute instances to use cases can reduce those costs, the same is true for storage. There are different storage needs for various OpenSearch Service use cases, and improving the alignment of the use case to the storage option will help reduce costs.

AWS is delivering new options, such as UltraWarm storage, to market. These offer different capabilities than S3. UltraWarm storage provides a cost-effective way to store large amounts of read-only data on OpenSearch Service. For example, log analytics customers can use a combination of Index State Management (ISM) for automating routine tasks and UltraWarm storage to reduce costs. Another option is



Cold Storage, which allows a customer to store large amounts of infrequently used or historical data on an OpenSearch Services domain at a lower cost. Cold Storage can be accessed on-demand for analysis or other uses. Using a combination of storage options, it is possible to reduce costs by as much as 82%. AWS also provides hourly backups for 14 days when you use OpenSearch Service, helping you save on backup costs.

In addition to lowering cost, OpenSearch Services storage options offer performance improvements. The latest generation of storage, GP3, not only reduces costs by 10%, but it also gives you the option to scale IOPS and throughput independently from storage. When coupled with Graviton, it is possible to increase storage density by 1.5x per instance. However, analysis done by AWS indicates that only 22% of customers are using GP3, making it an attractive option for immediate cost and efficiency improvement. ■

Chapter 4

Securing Critical Data in OpenSearch Service

Improving security has become a strategic task for most organizations. Protecting any solution that has sensitive or secret data is essential, and repositories, such as those that form the foundation of OpenSearch Service, are a cybersecurity priority. Given the large amounts of data, and the importance of the data in these systems, comprehensive security is an absolute must.

Security for OpenSearch Service is constantly enhanced to improve protection. This service offers a wide array of tools and technologies to provide advanced security options, similar to other AWS services. Key security capabilities include the following:

- Enhanced authentication and authorization.
- Encryption.
- Fine-grained access control.
- Access policies and network isolation.
- Audit logging and compliance.
- Secure programmatic access for custom applications and other AWS services.

AWS provides backward-compatible security updates for all major versions of OpenSearch and Elasticsearch, even those that are at end-of-life (EOL). Self-managed customers must upgrade immediately to get access to security patches or risk the security of their deployments, due to EOL. Using OpenSearch Service provides the flexibility to upgrade when the time is right, instead of adding in unplanned work when new Common Vulnerabilities and Exposures (CVE) are identified from an older version.

It is important to understand that upgrading to more current versions improves security. Upgrading to OpenSearch on the AWS service provides more security protection with encryption (in transit and at rest), SSL, HTTPS, fine-grained access control, credential/role-based authentication and authorization and more. There are also existing integrations with some of the most widely used third-party identity management systems.

Security for OpenSearch Service is constantly enhanced and improved to protect against the latest cyber threats. This service offers a wide array of tools and technologies to provide advanced security options.



AWS has added Amazon Cognito identity and access management into all versions of Elasticsearch from 5.1 forward. This increased level of security not only helps to meet an organization’s internal standards, but it also supports meeting regulatory and compliance demands if they apply. At present, this service supports multiple compliance programs, including HIPAA, FedRAMP, DoD CC SRG, SOC, PCI, ISO & CSA STAR, and FIPS 140-2.

AWS analysis shows that a full 80% of companies using the service can substantially improve their security posture simply by moving to the latest version and enabling the security options that become available for their existing domain.



There are some simple steps that customers can take to improve security that are available to them in the OpenSearch Service. There are two excellent examples of straightforward ways to improve security:

- **Implementing node-to-node encryption.** This security feature adds TLS 1.2 encryption for all communications within the virtual private cloud. The AWS analysis shows this feature should be utilized by more customers, as only 55% of customers are using it, despite availability from the Elasticsearch 6.0.
- **Deploying fine-grained access control.** This is a fundamental security function, providing a specific user/account with access to the minimal amount of information that they need. This helps reduce lateral spread. Implementation of this capability should be quite common, but the AWS analysis indicates that only 25% of customers have chosen to implement it.

The implementation of these security capabilities and the continual commitment to improving security in the OpenSearch Service make it possible to confidently host secure workloads and protect/limit access to confidential data. AWS’s seamless delivery and the updating of these features simplify or eliminate many tasks that previously burdened internal IT and security operations teams. ■

Chapter 5

Enhancing Availability and Redundancy

OpenSearch Service supports mission-critical workloads for both analytics and monitoring use cases. This makes it essential that the service is resilient and available. Unplanned downtime can create data integrity issues, interrupting any number of important workloads and processes.

There are very real and substantial differences in availability and resiliency for OpenSearch environments based on the way they are deployed and supported. OpenSearch Service has unique technology and supporting processes to provide better resilience than other options.

AWS is committed to highly resilient services, and customers that follow best practices and guidelines are guaranteed to have 99.9% availability. These guidelines include the use of automated failover, implementation of constant monitoring, load balancing and the effective use of AZs. AZs are distinct locations within an AWS Region that are engineered to be isolated from failures in other AZs. The use of multiple AZs is an important tool for improving the resilience of OpenSearch Service. There are numerous AZ options for customers, as the service is available in 29 Regions globally.

The service also makes it possible to have automated failover across AZs to provide uninterrupted services, handled by the managed service. AWS has a documented process for “shard” distribution from primary zones to secondary zones. This service also provides for automated and manual snapshots of the indexes and states. State parameters include cluster settings, node information, index settings and shard allocation.

Beyond this level of protection, AWS copies the dedicated master node to three AZs, even when the customer chooses to use only two. These zones are physically separate, reducing the impact of any natural disaster or other event that may impact a specific geographic area. The [AWS resiliency guidelines](#) call for the use of two, and ideally three, AZs for better fault tolerance and improved availability. Unfortunately, the AWS analysis of current customer usage shows that only 25% of customers are following the published best practices. The data also indicates that of all OpenSearch Service customers, only 25% are using two AZs, and 20% are using three. Using multiple AZs is a simple improvement that noticeably reduces any impact from a service interruption.



There are very real and substantial differences in availability and resiliency for OpenSearch environments based on the way they are deployed and supported.

To ensure that resiliency and availability are at the highest levels, OpenSearch Service provides easy monitoring and automated failure detection and recovery. This helps simplify operations for the IT and SecOps teams, as many manual tasks are now automated. In some cases, events impacting the service are remediated before any staff become involved. In addition, moving to the cloud eliminates service interruptions that are caused by legacy infrastructure outages.



One of the most common mistakes that impacts resiliency is the continuing use of older or obsolete versions of OpenSearch or Elasticsearch. The latest AWS features support improved resiliency and are

found only in the most up-to-date software. Still, the use of legacy versions of software remains commonplace. The AWS analysis shows that only 28% of customers are using the latest versions of OpenSearch and are availing themselves of the most advance features and innovations.

A Simple Path to Improved Availability/Resiliency

The imperative to improve resiliency and availability for OpenSearch Service should lead businesses to take a few important steps that will result in a noticeable improvement:

- Upgrade to the latest version of software.
- Implement the best practices and guidelines for improving availability for AWS services.
- Ensure features that improve resiliency and availability are deployed and being used effectively.

Most of these changes can be made at minimal cost and with little to no disruption in regular operations. As more mission-critical activities leverage OpenSearch Service, downtime and outages will have larger and more problematic impacts. The good news is that resiliency can be easily improved by taking the three steps listed above. ■

Chapter 6

Tying It All Together for Better Outcomes

The business-critical use cases and applications supported by OpenSearch Service make a compelling case for optimizing its use. OpenSearch Service is quickly becoming a foundational technology solution for many organizations, as it provides critical information that is necessary for a data-driven business. This service provides insights essential to business, IT and security operations teams.

The ability to search, analyze and manage critical data provides visibility into business operations that allows managers and executives to make the most accurate and impactful decisions. The ability of OpenSearch Service to provide data from multiple sources and systems is a primary reason for its increasing adoption and usage. With consistent and verifiable data that supports the search and analytics capabilities of the service, users have more confidence in their decision-making.

Leveraging AWS's cloud expertise and the development of specific features and capabilities designed specifically for this service dramatically improves the customer experience in several ways.

To begin, the service ties together the security, cost efficiency and resiliency requirements for a business-critical digital solution in a single service. AWS has developed specific guidelines for resiliency and availability, guaranteeing 99.9% uptime and providing seamless integration of security enhancements to the service.

OpenSearch Service makes it easy to quickly deploy this solution for organizations that want to optimize their use of OpenSearch. It is now possible to focus on utilizing the solution without spending time on system administration, managing security for these products or diverting scarce resources to the daily operational needs of infrastructure. Simplifying the operational processes for OpenSearch Service moves the focus from operations to getting answers that improve business results. ■

Leveraging AWS's cloud expertise and the development of specific features and capabilities designed specifically for this service dramatically improves the customer experience in several ways.



Chapter 7

Conclusion and Next Steps

The use of OpenSearch Service is increasingly widespread. The ability to search and manage non-traditional data sets, including logs, security events, documents, and analytics, makes it a great solution for many tasks that form the foundation of a digital business. Using an index to search, it is possible to query extremely large data sets very quickly.

AWS has invested to deliver a complete cloud service that enables organizations to more easily and effectively use the OpenSearch analytics and search solution. AWS offers a free consultation to help potential users understand the unique value and benefits for their organization. ■

For more information about OpenSearch Service, please go to: <https://aws.amazon.com/opensearch-service/>. To learn about how Autodesk uses OpenSearch Service to gain new insights into log analytics, go to: <https://aws.amazon.com/solutions/case-studies/autodesk-log-analytics/>. For an example of how this service improves cybersecurity for Pearson, check out: <https://aws.amazon.com/solutions/case-studies/pearson-opensearch-service-case-study/>.
