

AMAZON EBS + SIGNALFX INTEGRATION

What Is Amazon EBS?



Amazon's Elastic Block Store (EBS) is used to provide disk volumes for Elastic Compute Cloud (EC2) instances. It's elastic because it allows users to quickly create volumes based on their needed size and performance. EBS also offers greater durability by automatically replicating within an availability zone, making it easy to setup encryption and create snapshots.

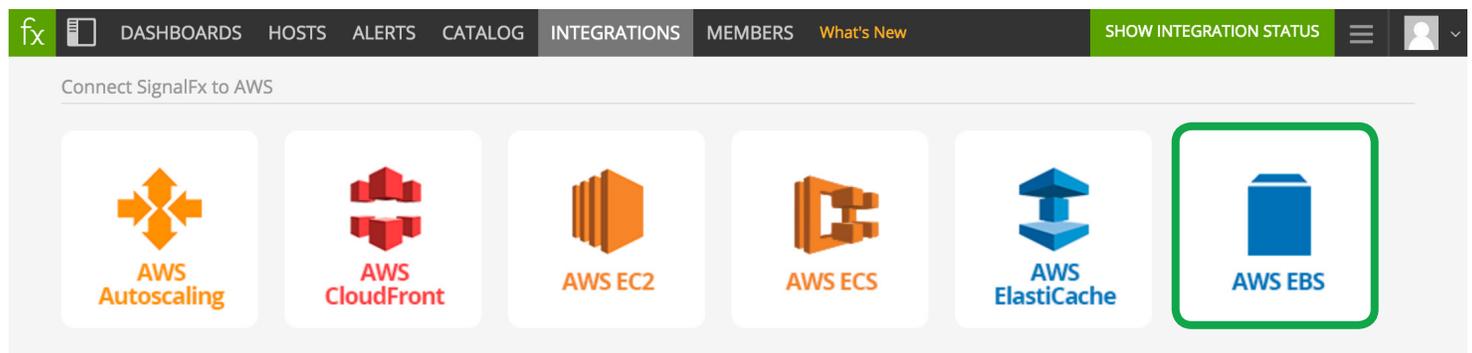
Monitoring Amazon EBS

Since disk volumes are a foundational resource for operating systems and their applications, it's important to monitor them to ensure your application's availability. If your disk has hit a performance limit, it can slow down your application. If the disk is full, it may even crash the server.

SignalFx offers a dashboard with pre-built content straight out of the box so that you can easily get insight into factors like read latency and throughput rate without having to set up your own trend graphs.

VOLUMES: Performance of EBS is determined by the type of volume used. There is a 500% difference in cost per GB between the least expensive hard disk drive (HDD) and the most expensive solid state drive (SSD), so monitoring relative to cost is essential. Monitor I/O per second (IOPS) and throughput during peak usage or stress test to determine if they are critical factors in your application performance. If your application is CPU- or memory-limited, you might be able to save money by switching to a lower-cost volume type.

TOTAL OPERATIONS: Seeing the total I/O operations and throughput across your entire cluster gives you an idea of how much capacity you're using in aggregate. Large spikes might indicate a system-wide activity pattern, like a sudden increase in traffic to your website. A good load balancer would distribute traffic, so you'd expect to see a fairly even distribution across your servers.



READ OPERATIONS: A breakdown of read operations both by the top volumes and by percentile distribution helps determine which volumes are using the most read ops. Additionally, a percentile time series chart shows how many servers are affected and how the load changes over time. Looking at the number of bytes per read operation also provides insight into how many volumes are experiencing small reads versus large reads. Small reads are better served by an SSD, which has more IOPS capacity. HDDs are a better fit for large reads because of their higher 1k maximum IO size (or block size).

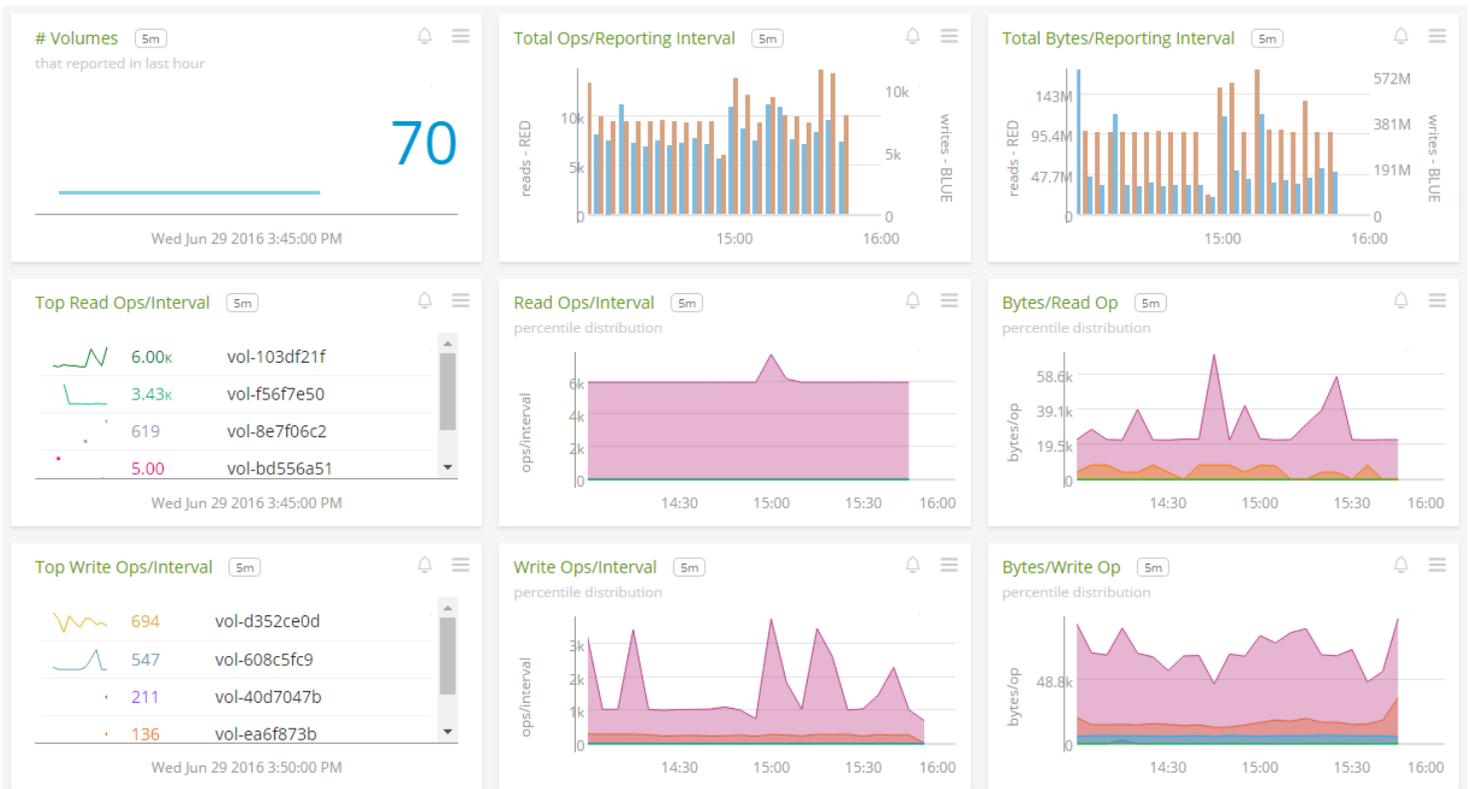
Amazon EBS Metrics

Volume Read Bytes	Volume Write Bytes
Volume Read Ops	Volume Write Ops
Volume Total Read Time	Volume Total Write Time
Volume Idle Time	Volume Queue Length
Volume Throughput Percentage	Volume Consumed Read Write Ops

WRITE OPERATIONS: Write ops are reported similarly to read ops. This helps you determine how many servers are experiencing a high number of write operations through the percentile distribution. Likewise, you can see the number of bytes per operation, which could help you provision the best type of volume.

TOP QUEUE LENGTH: How many read and write operations are waiting to be completed? If your application has a high number of IOPS, you might be able to decrease the latency by keeping a low queue length. A higher queue length allows you to achieve higher throughput due to sequential reads or writes but at potentially a lower latency. SSDs should have a queue length of 1 per 500 IOPS, and HDDs should have at least 4 while performing 1 MiB sequential I/Os.

READ & WRITE LATENCY: Here, latency is the end-to-end time to complete an I/O operation, reported as a percentile distribution, including statistics like average and maximum times. If latency is higher than expected, IOPS or throughput may have hit a limit or queue length may not be optimized for volume type.



Monitoring Amazon EBS with SignalFx

PER VOLUME METRICS: SignalFx also provides access to metrics on a per-volume basis, which are useful when drilling down to troubleshoot issues with a specific volume. The per-volume dashboard shows operations per minute, throughput, and latency. The charts give an operational view of both read vs. writes, as well as percentage change for each relative to 24 hours ago.



VOLUME SIZE & UTILIZATION: Surprisingly, Amazon CloudWatch does not offer metrics about volume size or percent utilization. It's important to monitor disk size because, if your server runs out of disk space, it can become unstable or crash. The SignalFx plugin for collectd (along with the df plugin) helps monitor disk usage and compute additional metadata, making it easier to analyze. Other collectd plugins provide insight into page file usage, disk operations by process, file counts, network usage, and more.

Percent change in disk space by day Detector

Time -1w



CLOSE

SAVE



Detector description

Signal:

Select a Signal to Test Against

SELECT A df_complex.used host:integration-test-couchbase-server-4.1-0 plugin_instance:opt-couchbase-var

1 ts

SELECT B A

SELECT C (A-B)/B*100

INTELLIGENT ALERTING: You can also easily create custom calculations for EBS automatically in the SignalFx interface without learning a new programming language. For example, you can create a dynamic threshold based on a percentile by applying analytics like time shift and mean with just a few simple clicks. An alert detector that notifies you via PagerDuty, Slack, or email when disk space decreases by 30% within any given one-day period is a more meaningful signal that will result in fewer false-positives than a basic static threshold.

About SignalFx

SignalFx is the most advanced monitoring and alerting solution for modern infrastructure. Our mission is to help cloud-ready organizations drive high levels of availability in today's elastic, agile, distributed environments. With SignalFx, development and operations teams gain a real-time view of, interact with, and take action on the infrastructure and application metrics that matter. We have enterprise customers including Yelp, Cisco, Zuora, and Hubspot and thousands of users analyzing billions of metrics every day. SignalFx was founded in 2013 by former Facebook and VMware executives, launched in 2015, and is backed by Andreessen Horowitz and Charles River Ventures.

