What is Amazon EC2?

AWS EC2

Amazon's Elastic Compute Cloud (EC2) is used by 84% of companies on AWS. It provides the foundation for many organizations' cloud strategies, enabling teams to allocate compute resources rapidly and easily meet demand at both high and low points for truly web-scale performance. However, despite Amazon EC2's resilience and elasticity, there are still ongoing objectives that require close monitoring of capacity, predictability, and interdependence with other services and infrastructure.

Limitations of Amazon CloudWatch

Although Amazon CloudWatch is a popular monitoring tool for AWS services by default, its limitations make it insufficient for most organizations trying to monitor EC2 at any kind of scale.

First, it only provides operating system-level metrics like CPU and memory—it does not offer insight into your application layer. Servers are often part of complex systems, and you'll want to correlate between operating system-level metrics and applications or between multiple applications.

Second, CloudWatch only gives you two weeks of retention for your metrics data. It's helpful to look back several weeks or months ago to see changes that happen over time and put patterns into perspective across deployments and system changes.

Third, CloudWatch only offers the ability to create simple dashboard widgets with a single metric or to set alarms with simple static thresholds. This is a good starting point for simple systems, but more complex systems and larger teams usually need more advanced analytical power. Analytics allows them to predict problems before they occur using calculated fields or to cut down on annoying false-positive alerts using dynamic thresholds.

Monitoring Amazon EC2

SignalFx provides real-time monitoring and intelligent alerting for all the data sources across your modern stack. It performs analytics on metrics as they stream from EC2, plus any custom metrics you designate, aggregated with

Amazon EC2 Metrics

CPU Credit Balance	CPU Credit Usage	Active Hosts
Active Hosts by Instance Type	Active Hosts by Availability Zone	CPU %
Top Instances by CPU %	Top Images by CPU %	Disk Ops/Min
Disk I/O Bytes/Min	Disk Metrics 24 H Growth %	Network Bytes In
Network Bytes Out	Top Network Bytes In/Min	Top Network Bytes Out/Min
Network Bytes In vs. 24 H Change %	Network Bytes Out vs. 24 H Change %	

metrics from the rest of your cloud infrastructure and services in your environment, with 13-month retention to see more changes over time. You also get a built-in Amazon EC2 monitoring dashboard and recommended detectors right out of the box so you can see the metrics that matter to performance without guesswork or painful trial-and-error.

RESOURCE STARVATION: Servers can become unavailable when the resources that they need to support clients are exhausted. For example, web servers can become unresponsive when they lack sufficient CPU or memory to respond before timing out. If a server does not even have enough memory to support an incoming SSH connection, you will not be able to access it through a remote terminal. In this case, you'll need to do a hard reboot, which risks losing the system state.

A good monitoring system will store metrics from the instance and can show you an increase in its resource usage until eventually hitting a ceiling and becoming unavailable. Memory, in particular, is one of the most constrained resources on servers, so it's important to have visibility when you are hitting a limit. Optimally, SignalFx automatically derives metrics in its built-in Amazon EC2 dashboard to determine your remaining capacity for each key resource in days, based on consumption trends and service patterns.

USAGE & PERFORMANCE: Performance impacts user experience and, therefore, impacts earnings. For web applications, slow page load times will lead users to abandon the page, some never to return. The performance of the server you're running on underlies and determines application performance. It's important to continuously monitor system performance because it is often most affected during bursts of activity or periods of peak demand. You'll also want to drill down into performance data across several dimensions to determine the root cause of problems and address bottlenecks.

- **Top Images by CPU %:** Narrowing down instances using the most CPU by image may help you identify the cause of the problem faster. If you recently upgraded to a new image, you can see if that image has a different performance pattern. Additionally, if your organization uses custom images based on service or application types, it will help you determine the root cause at an application or service level.

- **Disk Metrics 24 H Growth %:** It's useful to see the change in disk usage today versus the same time yesterday. It can help you determine if unexpected changes in disk usage are due to changes made over the past day, such as new code deployments or changes in user demand. If your application has a steady baseline usage, you may want to alert if it exceeds a certain threshold.

- **Top Network Bytes In & Out Per Minute:** It's important to see which instances are using the most bandwidth so you can determine if there's a problem on one of those instances. You might expect to find high bandwidth usage on a web server, but probably much less on an LDAP server, especially in a smaller company. In fact, high network bandwidth going to your LDAP server could indicate a security problem. Additionally, if you're looking to reduce your network usage, this gives you a good place to start.
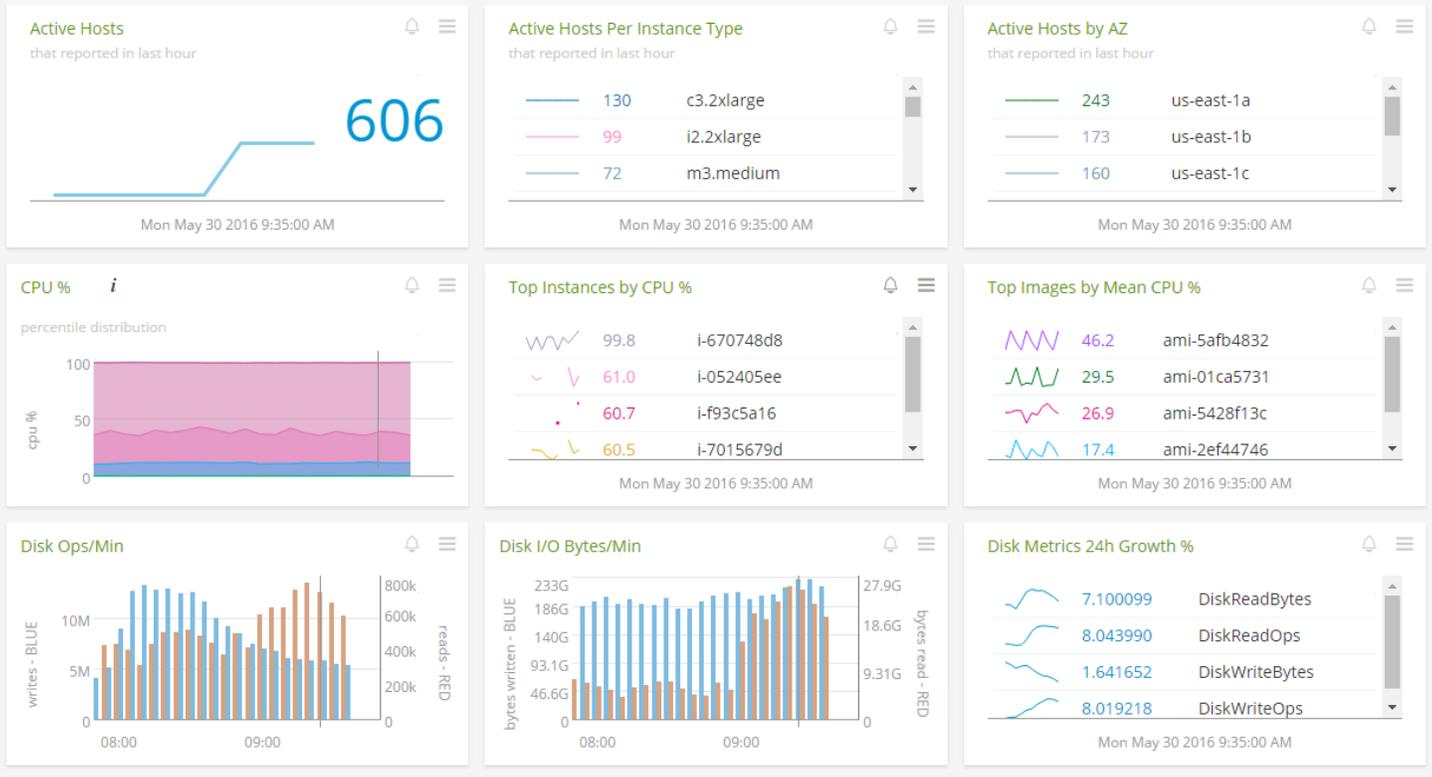
The SignalFx Difference

CPU BURST CREDITS: SignalFx also collects additional metrics from CloudWatch, including the number of CPU burst credits available. When you have spikes or temporary needs for CPU processing, you can use your burst credits to process the data quickly. If you need more CPU on a steady-state basis, you will need to increase your instance size. SignalFx also automatically tracks the number of network packets read and written over the network interface.

SignalFx's lightweight collectd agent is a valuable source of additional metrics for Amazon EC2. It offers a variety of plugins that can track memory usage, page faults, CPU steal time, disk space, process statistics, and other data points as time series, well beyond what you'd get from CloudWatch or other monitoring tools.

ANALYTICS & DERIVED METRICS: SignalFlow is the advanced analytics engine that allows you to take standard metrics and calculate new and more intelligent signals to monitor your systems and applications in real time as the data streams. For example, collectd's memory plugin lets you report memory usage either as an absolute number or as a percentage.

But what if you want both? Percentages are useful when you want to see whether a server is running below its maximum capacity or is at risk of running out of memory. For example, the ratio of memory used to total memory,

shown as percent memory used, provides a quick visual of the distribution of memory usage across all servers for the EC2 service and how it changes over time. If you know exactly how much memory each application uses, absolute numbers will help you determine how many more applications can fit on the server. You can also determine how much memory to add when upgrading the server.

By visualizing the data in multiple ways and combining time series as derived metrics with SignalFx, you can pinpoint the data streams that matter most and set dynamic alert thresholds from the built-in list of recommended detectors.
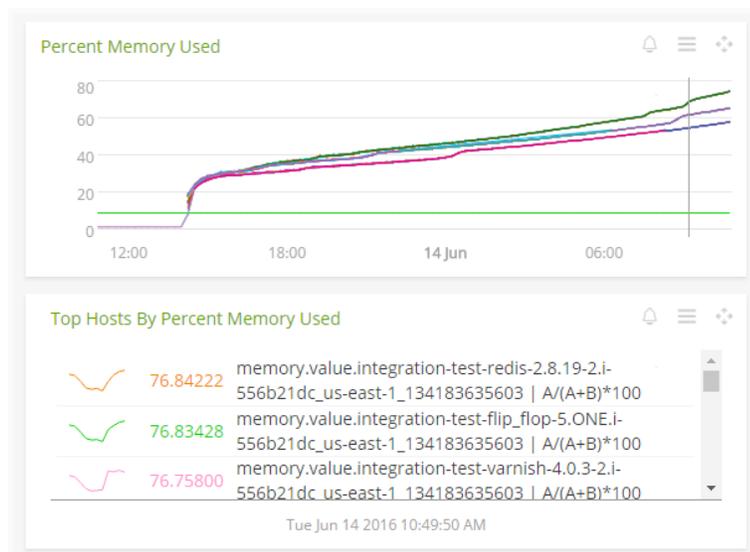
About SignalFx

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.