splunk> turn data into doing™

SOLUTION GUIDE

# Splunk Application Performance Monitoring (APM)

Implementation/launch resource

This resource is designed to provide implementation guidance to customers just getting started with Splunk® APM. Think of this as an easy-to-digest guide for implementations, training and services, and best practices.

Work through this entire guide at your own pace, applying these learnings to your APM environment as you go.

## Concepts and Value Proposition

In order to make the most of your investment in Splunk APM, it is important that you understand some core concepts and the value this product/technology aims to deliver.

### What is Splunk APM?

Splunk APM is an application performance monitoring and troubleshooting solution for microservices-based applications. APM monitors applications by collecting distributed traces. A trace is a collection of spans or actions that occur to complete a transaction.

### What problems does Splunk APM help solve?

Splunk APM helps with visualizing and understanding complex distributed environments that are critical to business functions, productivity and customer experience. It provides capabilities to reduce MTTR via unmatched levels of visualization and troubleshooting features.

### What are some business initiatives APM ties into?

- Adoption of the cloud and microservices — Trying to make the most of the benefits produced by the cloud, which includes developing and operating in the cloud along with microservices architectures
- Reduction of downtime — More engineering teams pushing more code, faster, always introduces risk of an outage or issue. APM provides the ability to better understand where an issue may be occurring in a distributed environment.
- Innovation — Confidence in operations and real-time visibility into the impact of changes foster innovation. Additionally, Lock-in with single vendors leads to higher prices without added value and slows down innovation. Development time spent integrating a proprietary single-vendor solution is wasted time.

### How does it work?

Through implementing instrumentation on the desired apps, APM collects and analyzes every span and trace that an application's instrumentation generates. This provides full-fidelity, infinite cardinality exploration of trace data an application generates, enabling you to break down and analyze application performance along any dimension.

### What does adoption of APM look like?

APM is all about answering the question "Where is the problem?" and providing guided troubleshooting workflows to pinpoint exactly where the problem may be occurring in a distributed environment. Splunk APM is powered by full-fidelity, no-sample tracing. This essentially means that Splunk APM looks at 100% of the data instead of a sample. For customers to adopt this product and gain the most value, they must send in trace data from their distributed environment via instrumentation; leverage this data for troubleshooting purposes when they are alerted of issues; and make use of the valuable features like tags + Tag Spotlight and Business Workflows.

**What are some outcomes we can expect from implementing Splunk APM?**

- Case study: One of the Year's First Unicorns Uses Splunk Observability to Conquer Cyber Five
- Case study: Care.com Refactors Monoliths Into Microservices With Splunk Observability

**Where can I go to get more help? — See Support Section of this document.**

- Open a Support Case — signalfx-support@splunk.com or within the Support Portal (which can be found in settings in-app)
- In-App Chat
- Documentation + More (Use sidebar and search to navigate) and Splunk Lantern

## Implementation

Splunk APM implementations consist of four high-level phases: Getting Data In + Training, Service Insights/ Views, Core Feature Configuration and Alerting.

### Getting data in + training

To begin, let's start with getting data in (GDI). In order to start sending data in, you will need to:

**Collect application data with an OpenTelemetry (OTel) Collector**

- As a first step to collecting data from your application, you should deploy the OTel Collector. This will allow you to export spans and traces from Kubernetes, Linux and Windows hosts and containers to Splunk Observability Cloud.
- To collect spans and traces from an infrastructure resource, select **Navigation menu > Data setup** and search for the host type or containerized environment you want to collect spans and traces from.
- Use the *environment* span tag to filter services by environment and easily monitor multiple environments separately.
- See these pages for more information about sending host or container data to Splunk Observability Cloud: Collect Kubernetes data, Collect Linux host data and Collect Windows host data

### Need extra help with Getting Data In?

There is an OnDemand Services task you can request to help unblock you:

- OTel Collector Configuration Guidance
- Smart Agent for Single Integration Configuration Guidance
- Assist with Auto Instrumentation

**Instrument your applications**

- Next, you can export spans to an OTel Collector running on the host or in the Kubernetes cluster that you deployed in the previous step. How you specify the OTel Collector endpoint depends on the language you are instrumenting. For more information, see the page for the language you are instrumenting in the list below.
- To collect spans and traces from a service, select **Navigation menu > Data setup** and search for an instrumentation library for the service you want to instrument.
- See the following pages to learn how to instrument a service or application running in each of these languages: Java, .Net, Node.JS, Python, Ruby and PHP

Once you have instrumented your applications, select **Observability > APM** and check that you can see your application data in the dashboard. If your data is not appearing in Splunk APM as you expect, see Troubleshoot your instrumentation.

Additionally, here is a list of all supported data sources and how to integrate them, for your reference. There's also an overview of important terms and concepts in Splunk APM.

**Recommended training to start:**

Using Splunk APM to Monitor Microservices-Based Applications

## Service insights/views

Now that we have our OTel Collector in place and our applications instrumented, we should now have data populating in the out-of-the-box visualizations in Splunk APM. Let's get a feel for these highly valuable components of the product.

### APM Homepage

The APM Homepage provides a high-density view at a service/workflow level with historical context. Out of the box, this page will show you Top Services by Error Rate, Top Services by Latency (P90) as well as the Top Business Workflows by Error Rate and Duration (P90). Choose your desired viewing preference for more details: Service Map, Tag Spotlight and Trace Search.

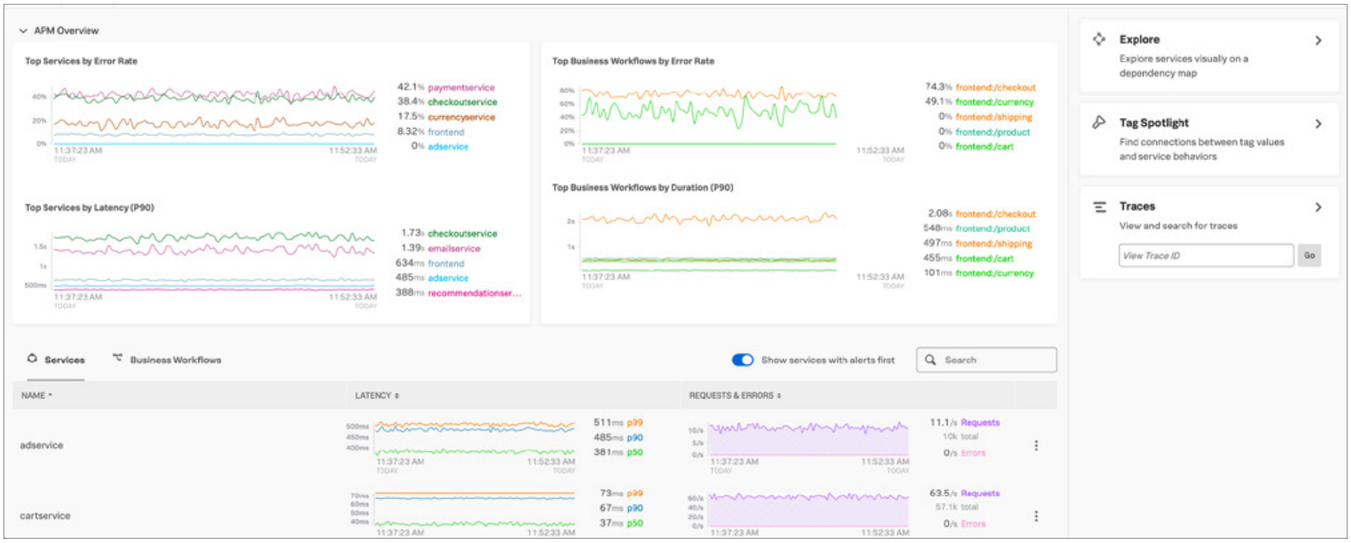

### Service Map

The Service Map is a visual representation of your various services and their dependencies. Splunk APM automatically discovers your instrumented services and their interactions to present dynamic and real-time service maps of your application's architecture. Use the service map to make more sense of your complex network of services and quickly identify the root cause at a glance, see latencies, dependencies, and slice and dice different services based on different tags.

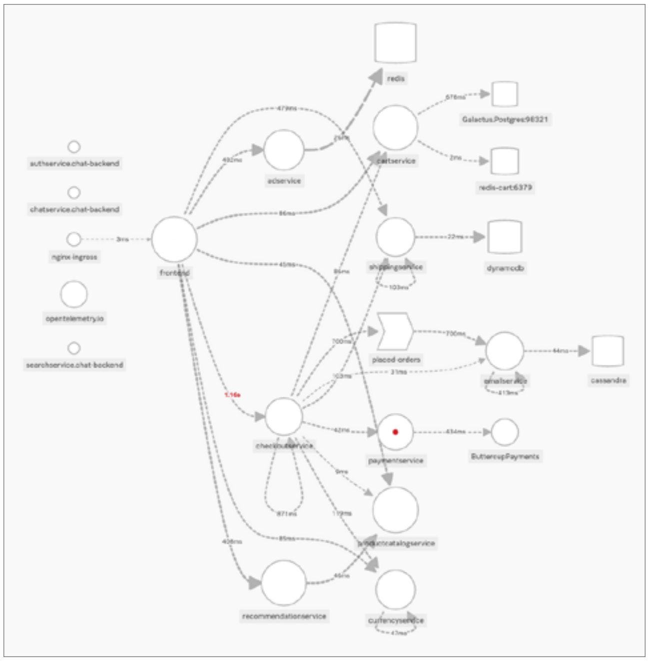

## Tag Spotlight

Use Tag Spotlight as the one-stop shop to analyze the performance of your services to discover trends that contribute to high latency or error rates with indexed span tags. You can break down every indexed span tag for a particular service to view metrics for it. When you select specific span tag values or a specific time range, you can view representative traces to learn more about an outlying incident.

For every service, Tag Spotlight provides a RED metrics time-series chart that displays the total number of requests, errors, root-cause errors and latency according to the specified time range in the APM navigation menu. Along with the RED metrics chart, Tag Spotlight also displays the total number of requests, errors, root-cause errors and latency for every value of an indexed span tag according to the specified time range in the APM navigation menu.

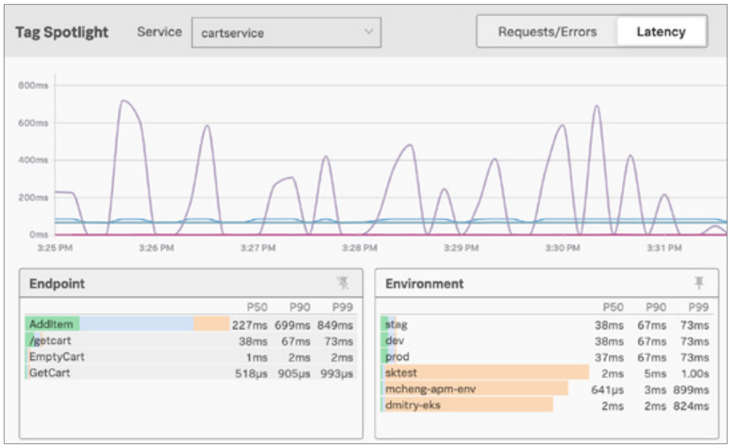

| Endpoint | P50 | P90 | P99 |
|---|---|---|---|
| AddItem | 227ms | 699ms | 849ms |
| /getcart | 38ms | 67ms | 73ms |
| EmptyCart | 1ms | 2ms | 2ms |
| GetCart | 518µs | 905µs | 993µs |

| Environment | P50 | P90 | P99 |
|---|---|---|---|
| stag | 38ms | 67ms | 73ms |
| dev | 38ms | 67ms | 73ms |
| prod | 37ms | 67ms | 73ms |
| sktest | 2ms | 5ms | 1.00s |
| mcheng-apm-env | 641µs | 3ms | 899ms |
| dmitry-eks | 2ms | 2ms | 824ms |

# High-value configuration items

There are a couple of high-value configuration items you will want to spend some time configuring and optimizing in order to get the absolute most out of Splunk APM. Indexing span tags will set you up to make the most out of Tag Spotlight, and configuring Business Workflows will unlock more seamless monitoring and troubleshooting of those critical flows throughout your distributed environments. Let's take a look.

## Indexing span tags

Drill down into service performance with span tags. Span tags provide additional context about operations that spans represent. Default span tags include things like the endpoint, operation and HTTP method associated with a span. Using these tags, you can analyze requests, errors and latency for spans that contain specific span tags. This context lets you understand service performance at a glance and helps you discover the root cause of issues faster.

Index span tags to analyze services in the following ways:

- Break down service performance by indexed tags in the Troubleshooting Service Map
- View charts of service performance metrics by indexed span tags in Tag Spotlight
- Track multiple traces for a specific activity with Business Workflows

### Which span tags to index?

Index only span tags you want to drill down into to gain insights about the performance of your infrastructure, or to address a specific incident. Some span tags provide a level of cardinality that just isn't useful. For example, indexing *query_id* would generate MetricSets for every unique query, and in most cases there's no reason for this level of cardinality. Also avoid indexing span tags that represent ephemeral resources, like *container_id*.

Consider which span tags are worth creating MetricSets for. Here are some questions you can ask about your environment:

- **Are there any attributes I look at when an incident occurs?** If you're running Kubernetes, you can index k8s.pod.name to view the performance of services by specific Kubernetes pods.
- **Do I run multiple versions or builds of code at the same time?** You can index tags for version or build_id to break down your infrastructure according to specific versions or builds of your applications.
- **Do I deploy services in multiple regions or fault domains?** It could be useful to view metrics for services by specific region span tags to identify issues with resources in specific regions or zones.

Here are the span tags that APM automatically indexes.

### How to index span tags

There are two ways to add span tags:

1. Instrument your application to create span tags.
2. Add span tags to spans when you send trace data to a Splunk OTel Collector.

**Instrument your application to create span tags:**

How you instrument code to create span tags depends on your code's language. For more information about adding span tags at the instrumentation level, see resources for the language you are instrumenting:

### Need extra help with Indexing Span Tags?

There is an OnDemand Services task you can request to help unblock you:

- Create Custom Span Tags

| Documentation | Instrumentation SDK |
|---|---|
| Instrument a Java Application | Splunk distribution of OpenTelemetry Java |
| Instrument a Node.js Application | SignalFx Tracing Library for JavaScript |
| Instrument a .NET Application | SignalFx Tracing Library for .NET |
| Instrument a Python Application | Splunk distribution of OpenTelemetry Python |
| Instrument a Ruby Application | SignalFx Tracing Library for Ruby |
| Instrument a PHP Application | SignalFx Tracing Library for PHP |

**Add span tags with an OTel Collector:**

Include span tags in settings for the batch processor in your OTel Collector configuration YAML file. You can create span tags with *attributes/newenvironment* which adds span tags to any spans that don't already have the tags or with *attributes/copyfromexistingkey*, which overrides an existing span tag value.

The settings look like this in an OpenTelemetry Collector configuration YAML file.

## Business Workflows

A Business Workflow is the start-to-finish journey of the collection of traces associated with a given activity or transaction. Each trace consists of multiple spans, and each span has identifying tags.

As a software engineer, site reliability engineer (SRE) or executive, you can use Business Workflows to monitor and troubleshoot end-to-end transactions in your system. In retail contexts, for example, an end-to-end transaction might encompass initial contact through order fulfillment, as captured by a trace.

You can create rules that correlate traces from a specific service or from multiple services that include the same global span tag. You must be an **administrator** to configure Business Workflow rules.

Check out this blog post going more into detail about how to improve business KPIs with business workflows.

### Configure a Business Workflow rule

To configure a new rule from Splunk APM, follow these steps. There is a difference between enabling a rule and applying it. The enable/disable switch affects an individual rule by turning it on or off. After you modify one or more rules, you then use buttons that act on the entire rule set to save or discard those changes. Changes are not applied unless you save them.

1. Go to **Organization Settings** (found at bottom of Nav Menu) **> Business Workflow Configuration**.
2. Click **New Rule**.
3. Select one of the following options from the **Rule Type** drop-down:
   a. **Global Tag** — Define workflows based on the value of a global tag in spans associated with a trace. This correlates traces that contain spans with the global tag.
   b. **Service** — Define workflows based on traces that include a service you specify. When a trace matches the rule, you also see a specified tag value or endpoint associated with the trace for the service.
4. Select a **Target Global Tag** or **Target Service** according to the **Rule Type** you selected.
   a. **Target Global Tag** prompts you to select an indexed global tag. When you select a tag, the rule correlates all traces with the global tag. The rule name is based on the global tag you select.
   b. **Target Service** prompts you to select a service and specify the **Source of Workflow Name**, which is extra metadata to view about the workflow. You can select to correlate traces for a service by an endpoint for the initiating span or a span tag value.
5. Click **Create** to save your changes and create the rule.
6. View the list of rules to confirm the rule you just created is enabled.
7. By default, the newest rule has the highest priority. This means Splunk APM applies the new rule before applying any other rules. If there are other rules you want to apply first, adjust the priority of the new rule.
8. Click **Save Changes** to apply the new rule and priority list.

Read more about configuring Business Workflow Rules here, and review an example rule configuration. You can also alert on Business Workflows which will be covered in the next section.

## Alerting

We've got our data flowing in, visualizations are populated and we're making use of Tag Spotlight as well as Business Workflows. What's next? Alerting, of course. APM detectors use built-in algorithms to detect sudden spikes and historical anomalies in your APM metrics or Business Workflows.

### Service and Business Workflow Detectors

You can dynamically monitor error rate and latency in the services and workflows you monitor with Splunk APM. Let's walk through a configuration of an APM Service/Business Workflow Detector below.

So, what can you configure within a detector? Detectors contain rules that specify:

- When the detector will be triggered based on conditions related to the detector's signal/metric.
- The severity of the alert to be generated by the detector.
- Where notifications should be sent.

From there, begin setting up your detector parameters:

- Type — Choose what type of detector to create: **APM Metric** or Infrastructure/Custom Metric (obviously selecting APM Metric in this case).
- Alert Signal — What **Service Metric** or **Business Workflow** are you trying to alert on? Your options: **Error Rate** or **Latency**. Here you will also define the specific **environment** and specific **service/endpoint**.
- Alert Condition — Define the conditions of the signal/metric in which you would like to be alerted on: Static Threshold or Sudden Change.
- Alert Settings — These settings will depend on which condition is selected and will be configured at this step.
- Alert Message — Define the severity of the alert and customize the message of it. You can also link to helpful documentation to be delivered with the alert.
- Alert Recipients — Define who will receive the alert and the delivery method: email, Splunk On-Call, Slack, PagerDuty, Webhook, etc.

### Need extra help with Alerting?

There is an OnDemand Services task you can request to help unblock you:

- Create a Simple Detector
- Create an Advanced Detector
- Detector Optimization

## Administration

Now that we have the core components all squared away, let's focus on administration. It's important for you to know how to best manage the tool in order to optimize usage throughout your organization.

You can find all of the documentation for administration related activities here, but let's touch on a few important ones to be aware of as you get started:

- Create and manage users — Add users to the instance and begin onboarding their data
- Create and manage access tokens — Use authentication tokens to authenticate Splunk API requests, track API usage and control your use of resources
- Manage permissions for detectors, dashboard groups and dashboards

## Training

Splunk offers a number of EDU training courses to help you get up to speed on how to make the most of APM. **Completion of these courses to some effect is an essential building block to success.** If you'd like to explore education options here, please get in touch with us via this contact form or get in touch with your account manager.

### Splunk APM course offerings and prices

- Using Splunk APM to Monitor Microservices-Based Applications — $500.00 USD or 1 Credit
  - This virtual, one-day course targeted at developers and DevOps enables you to use Splunk APM to analyze traces, troubleshoot and monitor your microservices-based applications. Through in-person discussions and hands-on activities, deep dive into uses of distributed tracing, navigating the Splunk APM app to analyze traces, visualize and alert on APM metrics.
- Advanced Monitoring of Microservices Applications Using Splunk APM — $500.00 USD or 1 Credit
  - This course, targeted at developers and DevOps, enables you to instrument your applications to send traces to Splunk APM. Through virtual discussions and hands-on activities, learn to deploy Splunk APM and use auto-instrumentation to send in traces without altering your code. Use manual instrumentation to create spans and add metadata to spans. You will also see how to configure and deploy the OTel Collector.

**Helpful links:**

Courses for Splunk Observability Customers

## Professional Services

Splunk's Experts are here to **partner with you to help achieve the outcomes that are important to your organization**. Access our experts through OnDemand Subscription Services or traditional Project-Based Services. We make it easy to get you the help you need in whatever way works for you.

### OnDemand Services

OnDemand Services (ODS) provides proactive technical adoption, implementation and optimization assistance for Splunk deployments, utilizing a pool of remote technical consultants. ODS credits are required in order to consume ODS, and the scope of ODS activities/tasks are predetermined. ODS can be requested within the Splunk Support Portal.

**How to request ODS:**

- Ensure your organization has ODS credits to use. If you do not, please get in touch with your Splunk point of contact to discuss how to purchase.
- Log in to the Splunk Support Portal.
- Use the navigation on the left to locate the OnDemand Services section and proceed with submitting your request.

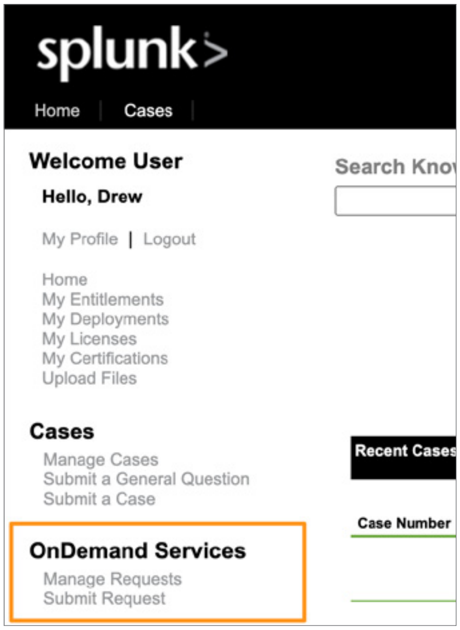

## ODS Catalog for Observability:

| Tasks: Observability Cloud, Infrastructure Monitoring (IM), APM, Log Observer (LO) | | | |
|---|---|---|---|
| **All Products:**<br>• Use Case Advisory Discussion<br>• Architecture Diagram Creation<br>**APM/IM/Cloud:**<br>• Cloud Migration Assessment | **APM/IM/Cloud:**<br>• Post Implementation Review<br>• Smart Agent for Single Integration Configuration Guidance<br>• OTel Collector Configuration Guidance<br>**Log Observer:**<br>• FluentD Configuration<br>• Log Processing Rule Configuration<br>• Metricization Rule Configuration<br>• Infinite Logging Configuration | **APM/IM/Cloud:**<br>• Create a Simple Detector<br>• Create an Advanced Detector<br>• Assist with Building a Simple Dashboard or Charts<br>• Assist with Building an Advanced Dashboard or Charts<br>**Cloud:**<br>• Getting Started with Splunk Observability Cloud<br>**IM:**<br>• Getting Started with Splunk Infrastructure Monitoring<br>• Assist with Exporting Data<br>• Assist with a Supported Cloud Integration<br>• Assist with a Supported Library Configuration<br>• Assist with the Configuration of prometheus-exporter<br>**APM:**<br>• Create Custom Span Tags<br>• Assist with Auto-instrumentation | **APM/IM/Cloud:**<br>• Usage Assessment<br>• Dashboard Administration Assistance<br>• Chart or Dashboard Optimization<br>• Detector Optimization |

**Helpful links:**

OnDemand Services Overview Video

## Project-Based Services

Project-Based Services are much more involved, typically larger-scale services engagements compared to ODS. With these, you will work with a Splunk Engagement Manager to determine and finalize the scope of the project. Once everything is signed off, we will work with you in lockstep to deliver on the agreed-upon project. If you'd like to explore options here, please get in touch with us via this contact form or get in touch with your account manager.

## Support

Even the most savvy customer will need a little help. Whether it's error messages, unexplained or unexpected behaviors, or incidents and outages, Technical Support is the first line of defense for all of your post-sales issues. Our Splunk Support Engineers will partner with you to ensure your environment is optimized to drive your journey with a focus on long-term technical health, so you can realize your ROI as soon as possible.

### How to open a support case:

- **Support portal:** This can be accessed from the application UI. Bring up the navigation menu, direct your attention to the bottom of the side-bar, select "Help & Support" and then select "Support and Community." From there you will be able to open a support case
- **Email:** Please email signalfx-support@splunk.com to open a support case
- **In-App chat** *(only available for customers with Premium Support entitlement by purchasing the Splunk Observability Cloud)*: A drawer or icon in the bottom right corner of the application is where you will find in-app chat. Engage there to be connected with a Support Engineer.

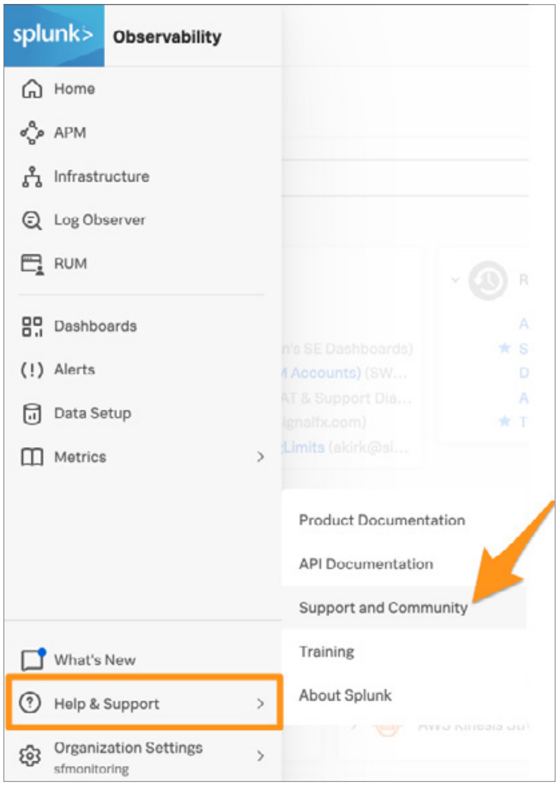

Get started with Splunk Observability Cloud today and reach out to us for on-demand expert help and implementation services.

Learn more: www.splunk.com/asksales

www.splunk.com