Splunk Enterprise on the AWS Cloud

Quick Start Reference Deployment

February 2017
Last update: May 2019 (revisions)

Bill Bartlett and Roy Arsan – Splunk, Inc.
Shivansh Singh – AWS Quick Start Reference Team

Contents

This Quick Start deployment guide was created by Amazon Web Services (AWS) in partnership with Splunk, Inc.

Quick Starts are automated reference deployments that use AWS CloudFormation templates to deploy key technologies on AWS, following AWS best practices.

Overview

This Quick Start reference deployment guide provides step-by-step instructions for deploying Splunk Enterprise on the AWS Cloud.

Splunk is a platform that makes machine data accessible and usable. By monitoring and analyzing everything from customer clickstreams and transactions to security events and network activity, Splunk software helps customers gain valuable Operational Intelligence from their machine-generated data. With a full range of powerful search, analysis, and visualization capabilities and prepackaged content for use cases, users can quickly discover and share insights.

Splunk Enterprise enables you to search, monitor, and analyze machine data from any source to gain valuable intelligence and insights across your entire organization. With Splunk Enterprise on the AWS Cloud, you gain the flexibility of the AWS infrastructure to tailor a deployment specific to your needs, and you can modify your Splunk deployment on demand, as these needs change. Lead times waiting for hardware to change or to scale your Splunk deployment are no longer a consideration with AWS.

This Quick Start is for IT infrastructure architects, administrators, and DevOps professionals who are planning to implement or extend their Splunk Enterprise deployments on the AWS Cloud.

Costs and Licenses

You are responsible for the cost of the AWS services used while running this Quick Start reference deployment. There is no additional cost for using the Quick Start.

The AWS CloudFormation template for this Quick Start includes configuration parameters that you can customize. Some of these settings, such as instance type, will affect the cost of deployment. For cost estimates, see the pricing pages for each AWS service you will be using.

This Quick Start requires a subscription to the Amazon Machine Image (AMI) for Splunk Enterprise, which is available from AWS Marketplace. For subscription instructions, see step 2 in the deployment steps. The AMI offers a 60-day trial license that provides limited access to Splunk Enterprise features. To fully utilize the environment created by this Quick Start, you will need to obtain a Splunk Enterprise license by contacting sales@splunk.com.

Architecture

Deploying this Quick Start **for a new virtual private cloud (VPC) with three Availability Zones** builds the following Splunk Enterprise environment in the AWS Cloud.

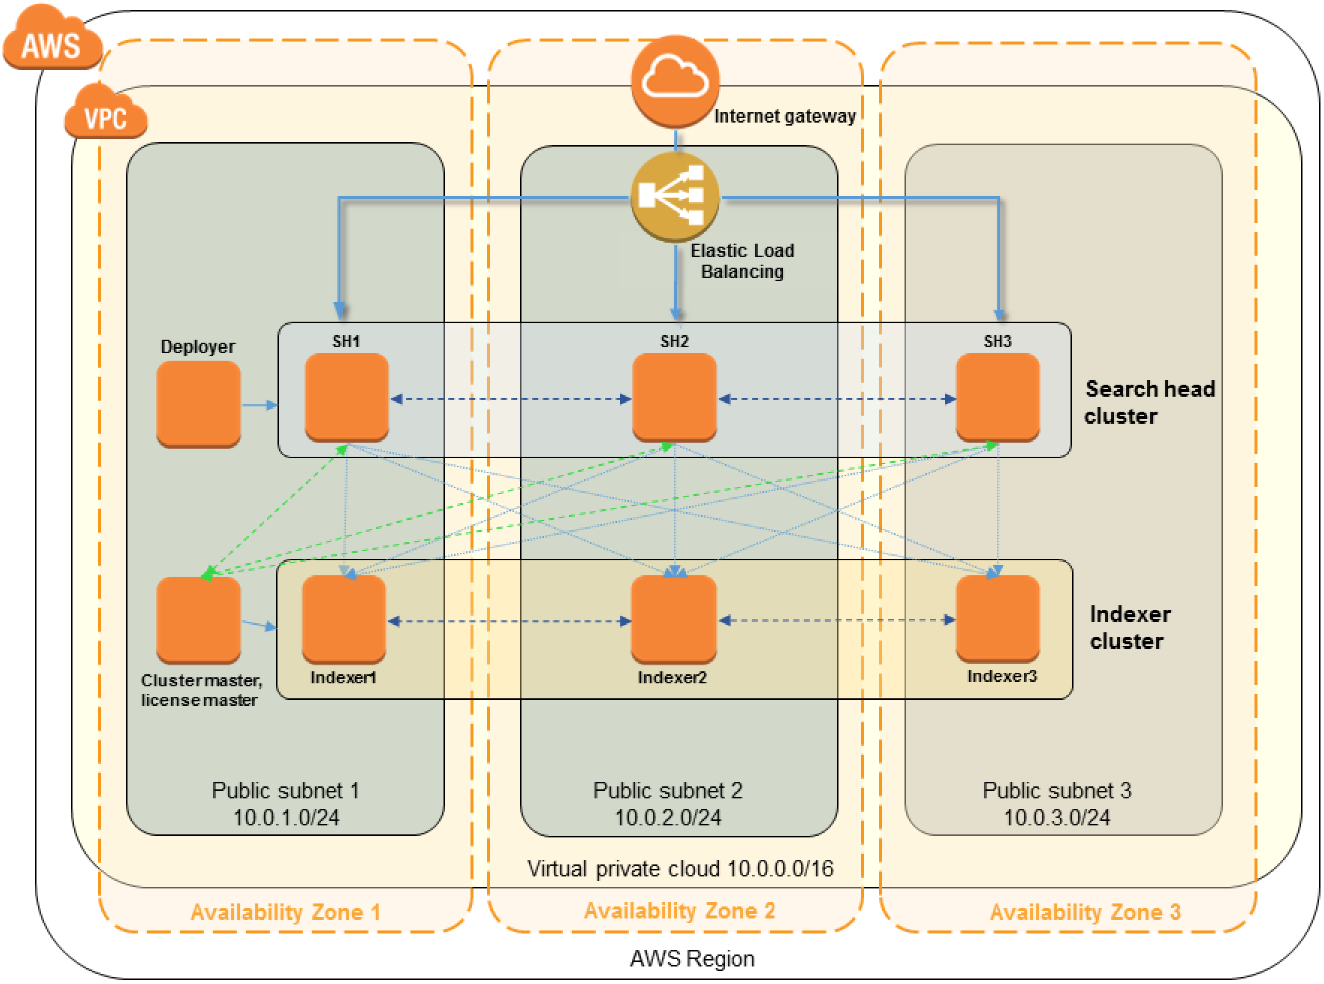

Figure 1: Quick Start architecture for Splunk Enterprise on AWS

The Quick Start sets up the following:

- A VPC configured across two or three Availability Zones. The Quick Start provisions one public subnet in each Availability Zone.
- Two Elastic Load Balancing (ELB) load balancers: one to load-balance HTTP web traffic to the search head instances, and the other to load-balance HTTP event traffic destined for the Splunk HTTP Event Collector (HEC) across all indexer instances.
- An IAM user with fine-grained permissions for access to AWS services necessary for the initial deployment process.
- Appropriate security groups for each instance or function to restrict access to only necessary protocols and ports.

- In the public subnets, EC2 instances for Splunk Enterprise, including the following:
  - Splunk indexer cluster with the number of indexers you specify (3-10), distributed across the number of Availability Zones you specify. The Splunk receiver (**splunktcp**) and Splunk HEC are enabled across all indexers.
  - Splunk search heads, either stand-alone or in a cluster, based on your input during deployment. In the latter case, the search heads are distributed across the number of Availability Zones you specify.
  - Splunk license server and indexer cluster master, co-located.
  - Splunk search head deployer, where applicable.
  - (Optional) User-provided Splunk apps and/or add-ons, loaded and pre-installed across indexers and search heads, based on your input.

If you decide to deploy Splunk Enterprise into your existing VPC (see Deployment Options later in this guide), the Quick Start assumes that the infrastructure components already exist, and deploys Splunk Enterprise into the environment you specify during deployment.

Prerequisites

Specialized Knowledge

Before you deploy this Quick Start, we recommend that you become familiar with the following AWS services. (If you are new to AWS, see Getting Started with AWS.)

- Amazon Virtual Private Cloud (Amazon VPC)
- Amazon Elastic Compute Cloud (Amazon EC2)
- Amazon Elastic Block Store (Amazon EBS)
- Elastic Load Balancing

You should also be familiar with core Splunk concepts, including both indexer clustering and search head clusters.

Planning the Deployment

Before you deploy Splunk Enterprise on AWS, please review the following sections for guidelines on nodes, instance types, storage, and high availability / disaster recovery (HA/DR) considerations for deployment.

Deployment Options

This Quick Start provides two deployment options:

- **Deploy Splunk Enterprise into a new VPC** (end-to-end deployment). This option builds a new AWS environment consisting of the VPC, subnets, security groups, and other infrastructure components, and then deploys Splunk Enterprise into this new VPC.
- **Deploy Splunk Enterprise into an existing VPC**. This option provisions Splunk Enterprise in your existing AWS infrastructure.

The Quick Start also lets you configure additional settings such as CIDR blocks, instance types, and Splunk Enterprise settings, as discussed later in this guide.

Single Instance vs. Distributed Deployment

Before you deploy a Splunk cluster, you'll need to decide whether to deploy a single node or create a distributed environment.

- A **single-node deployment** is one server providing all Splunk-related functionality, including license server, indexer, and search head. Typical use cases for a single-server deployment are small-scale, non-HA, low-volume production scenarios, proof of concept deployments, or dev/test/QA scenarios. We recommend a maximum of 300 GiB/day on a single-node deployment.
- A **distributed deployment** consists of several instances working together to provide indexing and search head duties. These deployments can range anywhere from two to hundreds of instances.

> **Note** This Quick Start provides a distributed (clustered) deployment to support high availability and higher volume environments. If you're interested in implementing a single-node deployment for the use cases discussed in this section, see the Splunk Enterprise on AWS manual deployment guide.

For further information about the dimensions of a Splunk Enterprise deployment on AWS, see the Splunk Enterprise capacity planning manual.

Storage

In most circumstances, the deployment type and use cases will dictate the type of storage you use.

When deploying a **non-clustered environment**, either single-server or distributed, we recommend utilizing EBS volumes and EBS-optimized instance types.

When architecting **a clustered solution**, two storage options are available:

- **Instance storage** is temporary, and is often referred to as *ephemeral storage*. If the instance is terminated or crashes, all data stored on the instance is lost.
- **An EBS volume** is persistent, even in the case of instance termination or crash.

Each option comes with its own feature set and price points. This Quick Start uses EBS volumes to help ensure high availability and durability for your instances.

Data Acquisition

You'll also need to determine how to best get your data into the Splunk platform.

- In most scenarios, we recommend using a **Splunk Universal Forwarder (UF)** installed on the machine where the data resides. A simple example of this would be installing the UF on a webserver to forward the webserver logs to the Splunk indexers.
- A **Splunk Heavy Forwarder (HF)** adds capabilities at the forwarding tier. In more advanced scenarios, an HF may provide the deployment with additional capabilities that the UF cannot.

For more information about these forwarders, see the Splunk Enterprise documentation.

If a forwarder can't be installed on the client that is generating data for the Splunk platform, sending that data directly to Splunk HEC is an option. In such situations, we recommend placing Splunk HTTP event collector(s) behind an ELB load balancer, and configuring the client to send the events to the load balancer address. This makes future scaling easier.

Splunk provides an add-on that automates data ingestion from several AWS services, including AWS Config and AWS Config rules, AWS CloudTrail, Amazon CloudWatch, Amazon Inspector, and Amazon CloudFront. Installing and configuring the add-on enables data collection from any of the services a user selects, directly to the user's Splunk Enterprise deployment.

This guide provides instructions for forwarding data in step 4 of the deployment instructions.

Instance Selection

As a general rule, deployments without Splunk premium solutions, and single-node (non-clustered) deployments should use C5 instance types. (For information on instance types, see the AWS website.) This Quick Start uses the c5.4xlarge instance type by default, but you can choose other C5 sizes as well as C4, M4, and I3 instance types. Which instance type to use depends on how much workload will be delivered to the instance. For Splunk Enterprise, the workload is defined as both indexing and search. In the following tables, recommendations are based on typical (average) use, but heavy searching can impact performance as much as, or more than, heavy indexing.

The assumptions in the following tables assume Amazon EBS General Purpose SSD (gp2) volumes attached to the C5 instance types. In all situations, we recommend deploying on dedicated hosts to avoid potentially noisy neighbor situations.

Indexers

Instance type	Daily indexing volume (GiB)
c4.2xlarge	<100
c4.4xlarge	100-200
c4.8xlarge	200-300+

Search Heads

Instance type	Concurrent users
c4.4xlarge	Up to 8
c4.8xlarge	Up to 16

When using Splunk premium solutions such as Splunk Enterprise Security (ES) or Splunk IT Service Intelligence (ITSI), we recommend indexer instance types that have a larger memory footprint. The following information also assumes Amazon EBS General Purpose SSD (gp2) volumes attached to the M4 instance types.

Indexers (Splunk Premium Solutions)

Instance type	Daily indexing volume (GiB)
i3.8xlarge	100
m4.10xlarge	100-150

Search Heads (Splunk Premium Solutions)

Instance type	Concurrent users
r4.8xlarge	Up to 16
m4.10xlarge	Up to 20

HA/DR Considerations

For deployments that require either high availability (HA) or resilient disaster recovery (DR) capabilities, Splunk Enterprise has built-in clustering technology to safely replicate data. There are two clustering options—search head clustering and indexer clustering—that allow for considerable flexibility when architecting a deployment.

Indexers

The Quick Start offers you a choice of 3 to 10 indexers to launch during deployment. (The default is 3.)

Indexer Replication

The importance of ensuring resiliency and recoverability of data in the event of a node failure or other disaster cannot be overstated. One way of doing that at the indexer tier is to use Splunk indexer replication.

Consider the following three questions when deciding how to provision indexer replication:

- Is a secondary site necessary? Does the deployment require a secondary physical location to mitigate losses in the case of a natural disaster or otherwise catastrophic damage to the primary location? This is called multisite clustering.
- How many copies of data should the cluster replicate? Essentially, how many failed nodes should the deployment be able to tolerate before data loss? This is called replication factor (RF). The number of concurrently tolerated failed nodes is RF–1. For example, if the replication factor is set to 3, then two nodes can fail without data loss. The Quick Start lets you choose 2-4 copies of data to replicate, and uses a default of 2 copies.
- How many immediately searchable copies of data should the cluster retain? This is called search factor (SF). It is possible to have a replication factor that is higher than a search factor if data resiliency is the primary requirement, as opposed to high availability search capability. The Quick Start lets you choose 2-4 copies of data to retain for searching, and uses a default of 2 copies.

In the case of multisite clustering, each distinct location is referred to as a *site* with the RF and SF configured on a per-site basis. For example, a primary and DR configuration might configure the total RF and SF to 4. The primary site could maintain three replicated and searchable copies while the backup/DR site could maintain one copy. A cluster cannot have a search factor that is greater than the replication factor.

In AWS deployments, a Splunk Enterprise site will typically be in an Availability Zone or an AWS Region, depending on requirements. Each Availability Zone is a physically unique facility, and many customers will find multisite deployments across multiple Availability Zones in a single region sufficient for HA. If your requirements are for redundancy across different geographic locations, treating an entire region as a site will fulfill those requirements.

This Quick Start deploys Splunk Enterprise into two or three Availability Zones (depending on your input) within a single AWS Region with multisite clustering, where each Availability Zone is considered a separate site. Both RF and SF are set to 2, so each site maintains a complete, searchable copy of the data. This helps ensure high availability of the indexer or data tier, and tolerates failure of one Availability Zone by default.

Figure 2 illustrates an indexer cluster where both SF and RF equal 3.

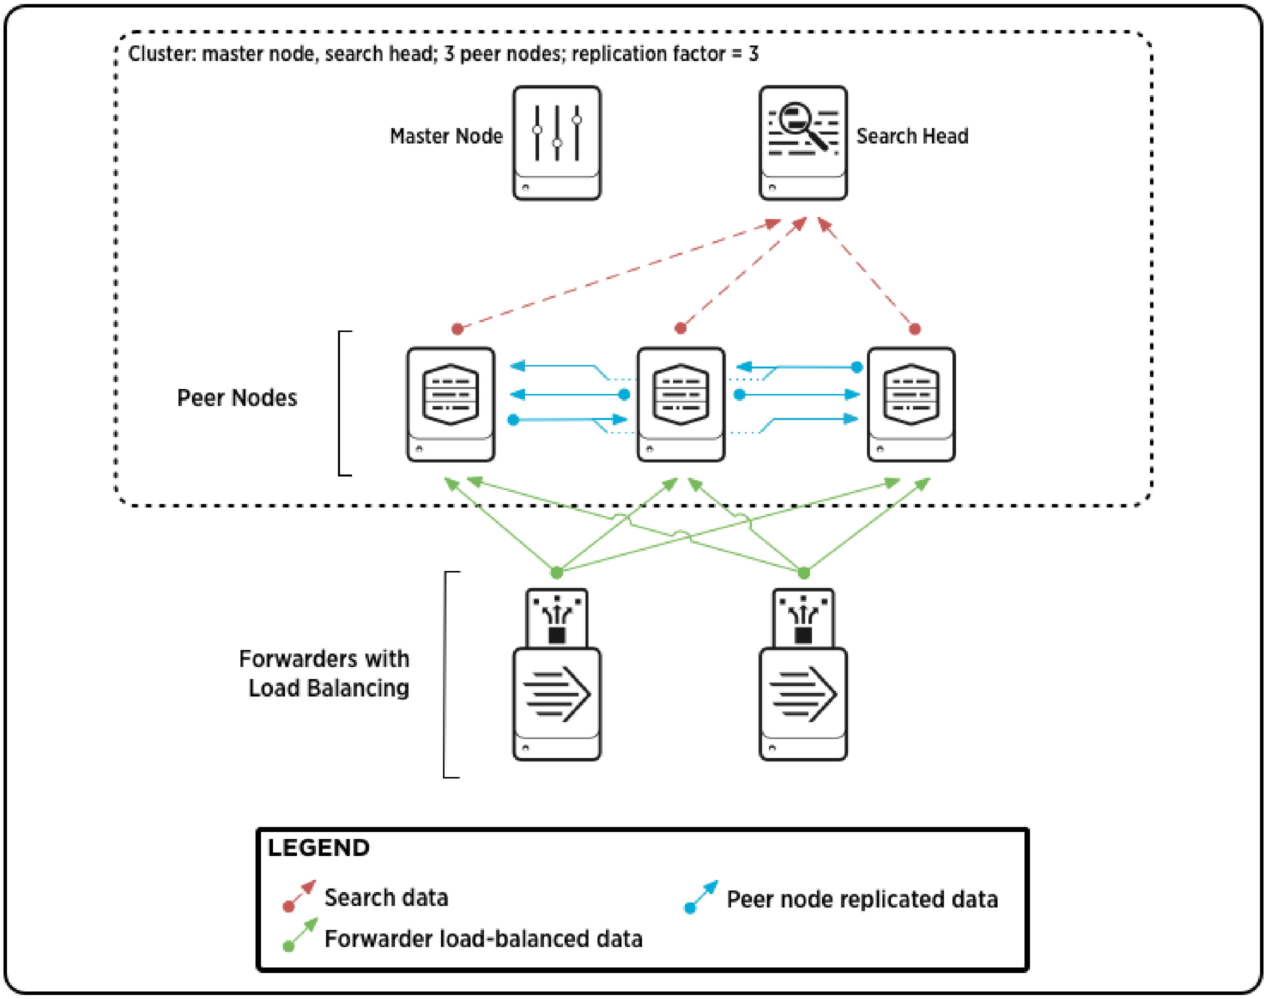
Cluster: master node, search head; 3 peer nodes; replication factor = 3
Master Node
Search Head
Peer Nodes
Forwarders with Load Balancing
LEGEND
Search data
Forwarder load-balanced data
Peer node replicated data

Figure 2: Splunk single-site cluster architecture

When deciding how to properly provision a clustered deployment, it's important to consider the differences in requirements for a replicated-only bucket versus a replicated and searchable bucket. For a replicated-only bucket, we estimate that roughly 15% of the original size of indexed data will be copied to the replicated destination(s). For a replicated and searchable bucket, the estimate increases to about 50% of the original indexed data. Lastly, note that a Splunk cluster will try to always maintain the appropriate number of replicated buckets. When a node fails, Splunk will try to re-replicate all buckets on that failed node across the cluster to maintain the configured replication count.

We recommend reserving, at minimum, one indexer's worth of storage across the cluster, to allow for such failures without filling up the storage volumes. For example, if a cluster has 10 indexers, and each indexer has 10 TiB of disk space, reserving a minimum of 10 TiB in total across the cluster (approximately 1 TiB per indexer) would be the suggested course of action. We recommend keeping a cluster around 80% capacity to allow for both failures and unexpected spikes in traffic.

Search Heads

Search Head Clustering

Similar to indexers, Splunk search head clusters allow for replication of configuration, job scheduling, and saved search results. The replication factor (RF) works the same way as for the indexers; a search head cluster can tolerate simultaneous failures from RF–1 nodes. Again, similar disk space considerations should be implemented to accommodate for artifacts being redistributed upon node failure. If you'd like to deploy a search head cluster, set the **SHCEnabled** parameter to **yes** during Quick Start deployment. With that setting, the Quick Start will deploy a 3-node search head cluster across the Availability Zones and will replicate artifacts 2 times (RF=2). This helps ensure high availability of the search tier, and tolerates failure of one Availability Zone by default.

Elastic Load Balancing

The Elastic Load Balancing service can provide a Splunk Enterprise deployment with an additional layer of fault tolerance in a few different scenarios. Deploying in front of a search head cluster is one such scenario.

When deploying a Splunk search head cluster, you should place all search heads behind a Classic Load Balancer or an Application Load Balancer to help ensure the highest availability. When you place the search head instances in different Availability Zones, and the load balancer directs traffic across only healthy search heads, any outage in a single Availability Zone will automatically be routed around by the load balancer. Similarly, when deploying distributed HEC input, you should place it behind a Network Load Balancer or an Application Load Balancer to ensure high availability and uniform load distribution.

This Quick Start automatically sets up the Elastic Load Balancing service, which provides HTTP load balancing across the search head instances for the web traffic, and HTTP load balancing across indexers for the data traffic destined for the HTTP Event Collector, which is enabled across all indexers. We highly recommend deploying your own SSL certificates and replacing the HTTP listeners with their HTTPS counterparts. For detailed steps, see the AWS documentation.

In contrast, the standard Splunk receiver ports across indexers accept traffic directly from Splunk universal forwarders, which automatically handle load balancing among other fault tolerance capabilities.

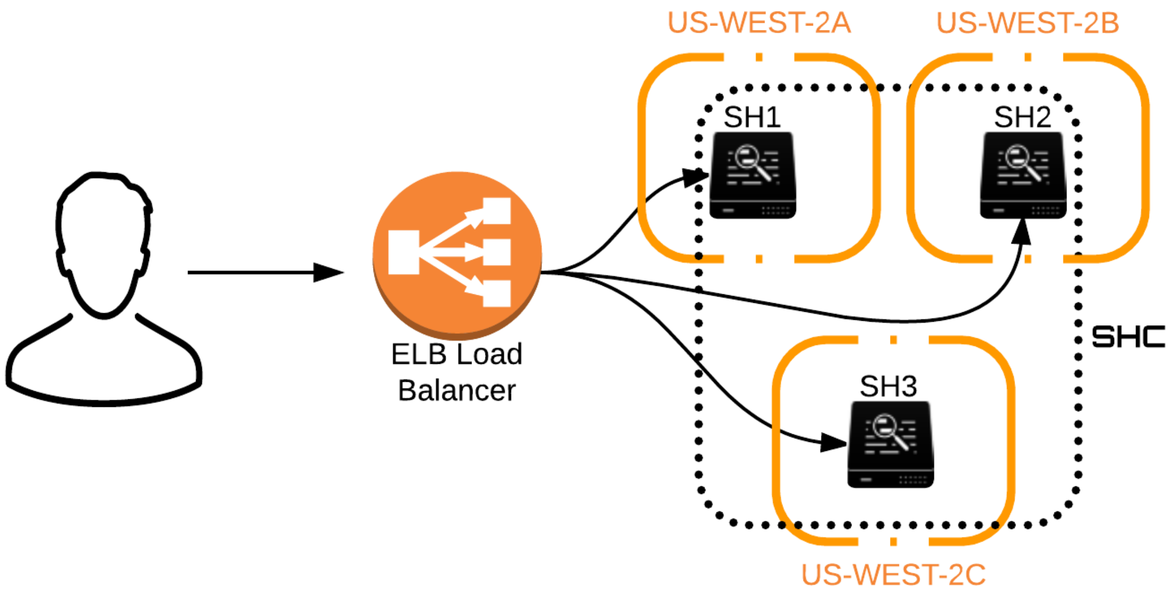

Figure 3: Elastic Load Balancing for clustered search heads

Deployment Steps

Step 1. Prepare Your AWS Account

1. If you don't already have an AWS account, create one at https://aws.amazon.com by following the on-screen instructions.
2. Use the region selector in the navigation bar to choose the AWS Region where you want to deploy Splunk Enterprise on AWS.
3. Create a key pair in your preferred region.
4. If necessary, request a quota limit increase for the EC2 instance type that you've decided to deploy Splunk Enterprise on (c4.large by default). You might need to do this if you already have an existing deployment that uses this instance type, and you think you might exceed the default quota with this reference deployment.

Step 2. Subscribe to the Splunk Enterprise AMI

This Quick Start requires a subscription to the Amazon Machine Image (AMI) for Splunk Enterprise running on Amazon Linux.

The AMI provides a 60-day free Enterprise trial license, which supports a limited set of features. To take full advantage of the Splunk Enterprise feature set, including distributed search, you can obtain a license for Splunk Enterprise by contacting sales@splunk.com.

To subscribe:

1. Log in to your AWS account.
2. Open the AWS Marketplace page for Splunk Enterprise, and choose **Continue**.
3. Use the **Manual Launch** option to launch the AMI into your account on Amazon EC2. This involves accepting the terms of the license agreement and receiving confirmation email. For detailed instructions, see the AWS Marketplace documentation.
4. If you're using a BYOL license, place the Splunk license key file in a private S3 bucket. You'll be able to enter the bucket name and the file path as part of the Quick Start parameters during deployment, as described in the next step.

Step 3. Launch the Quick Start

> **Note** You are responsible for the cost of the AWS services used while running this Quick Start reference deployment. There is no additional cost for using this Quick Start. For full details, see the pricing pages for each AWS service you will be using in this Quick Start.

1. Choose one of the following options to launch the AWS CloudFormation template into your AWS account. For help choosing an option, see deployment options earlier in this guide.

Option 1	**Option 2**
Deploy Splunk Enterprise into a new VPC on AWS	**Deploy Splunk Enterprise into an existing VPC**
Launch	Launch

> **Important** If you're deploying Splunk Enterprise into an existing VPC, make sure that your VPC has sufficient public subnets in different Availability Zones for the Splunk indexers and search heads, consistent with the number of Availability Zones you specify during deployment. You'll be prompted for your VPC settings when you launch the Quick Start.

   Each deployment takes about 10-30 minutes to complete, depending on whether you decide to enable search head clustering.

2. Check the region that's displayed in the upper-right corner of the navigation bar, and change it if necessary. This is where the network infrastructure for Splunk Enterprise will be built. The template is launched in the US West (Oregon) Region by default.

3. On the **Select Template** page, keep the default setting for the template URL, and then choose **Next**.

4. On the **Specify Details** page, change the stack name if needed. Review the parameters for the template. Provide values for the parameters that require input. For all other parameters, review the default settings and customize them as necessary. When you finish reviewing and customizing the parameters, choose **Next**.

   In the following tables, parameters are listed by category and described separately for the two deployment options:

- – Parameters for deploying Splunk Enterprise into a new VPC
- – Parameters for deploying Splunk Enterprise into an existing VPC

Option 1: Parameters for deploying Splunk Enterprise into a new VPC

View template

AWS Instance and Network Settings:

Parameter label (name)	Default	Description
EC2 instance type for Splunk indexer (IndexerInstanceType)	c5.4xlarge	EC2 instance type to use for Splunk Enterprise indexers. For guidance on selecting an instance type, see Instance Selection earlier in this guide.
EC2 instance type for Splunk search head (SearchHeadInstance Type)	c5.4xlarge	EC2 instance type to use for Splunk Enterprise search heads. For guidance on selecting an instance type, see Instance Selection earlier in this guide.
Key Name (KeyName)	*Requires input*	Public/private key pair, which allows you to connect securely to your instance after it launches. When you created an AWS account, this is the key pair you created in your preferred region.
Permitted CIDR for Splunk web interface (WebClientLocation)	*Requires input*	The CIDR IP range that is permitted to access the Splunk servers' web interface. We recommend that you set this value to a trusted IP range.
Permitted CIDR for Splunk HTTP event collector input (HECClientLocation)	*Requires input*	The CIDR IP range that is permitted to access Splunk HTTP Event Collector (HEC). We recommend that you set this value to a trusted IP range.
Permitted CIDR for SSH (SSHClientLocation)	*Requires input*	The CIDR IP range that is permitted to access Splunk Enterprise instances via SSH. We recommend that you set this value to a trusted IP range. For example, you might want to grant only your corporate network access to the Splunk deployment.
Availability Zones (AvailabilityZones)	*Requires input*	The list of Availability Zones in the AWS Region where you want to install Splunk Enterprise. The number of selections must match the value you specify in the **Number of Availability Zones** parameter. The Quick Start preserves the logical order you specify.
Number of Availability Zones (NumberOfAZs)	2	The number of Availability Zones to use in the VPC. This must match your selections in the Availability Zones parameter. You can choose 2 or 3 Availability Zones.
VPC CIDR (VPCCIDR)	10.0.0.0/16	CIDR block for the VPC.

Parameter label (name)	Default	Description
Public Subnet 1 CIDR (PublicSubnet1CIDR)	10.0.1.0/24	CIDR block for the public subnet located in Availability Zone 1.
Public Subnet 2 CIDR (PublicSubnet2CIDR)	10.0.2.0/24	CIDR block for the public subnet located in Availability Zone 2.
Public Subnet 3 CIDR (PublicSubnet3CIDR)	10.0.3.0/24	CIDR block for the public subnet located in Availability Zone 3.

Splunk Settings:

Parameter label (name)	Default	Description
Splunk Admin Password (SplunkAdminPassword)	*Requires input*	The password for Splunk Enterprise. This string must be 6-32 characters long, and may contain letters, numbers, and symbols.
Shared Security Key for Cluster Nodes (SplunkClusterSecret)	*Requires input*	Shared cluster secret for the Splunk search head and indexer clusters. This string must be 8-32 characters long, and may contain letters, numbers, and symbols.
Shared Security Key for Forwarders using Indexer Discovery (SplunkIndexerDiscovery Secret)	*Requires input*	Security key used for communications between forwarders and the cluster master. This value should also be used by forwarders to retrieve a list of available peer nodes from the cluster master. This string must be 8-32 characters long, and may contain letters, numbers, and symbols.
Splunk License Bucket (SplunkLicenseBucket)	—	The name of the private S3 bucket that contains your Splunk license key file, from step 2.
Splunk License S3 Bucket Path (SplunkLicensePath)	—	The path to the S3 bucket that contains your Splunk license key file, without a leading forward slash (/), from step 2.
No. of Splunk Indexers (SplunkIndexerCount)	3	The number of Splunk Enterprise instances to launch. You can choose from 3 to 10 instances.
Indexer Disk Size (SplunkIndexerDiskSize)	320	The size of the EBS volume attached to the Splunk Enterprise indexers, in GiB. You can choose a value between 320-16,000. For guidance, see Instance Selection earlier in this guide.
Search Head(s) Disk Size (SplunkSearchHeadDisk Size)	320	The size of the EBS volume attached to the Splunk search head, in GiB. You can choose a value between 320-16,000. For guidance, see Instance Selection earlier in this guide.
Index Cluster Replication Factor (SplunkReplicationFactor)	2	The number of copies of data to store in the Splunk indexer cluster. You can choose 2-4 copies. For guidance, see Indexers earlier in this guide.
Index Cluster Search Factor (SplunkSearchFactor)	2	The number of copies of data to retain for searching in the Splunk head cluster. You can choose 2-4 copies.

Parameter label (name)	Default	Description
Enable Search Head Cluster? (SHCEnabled)	no	Set this parameter to **yes** to deploy a Splunk search head cluster. (The default setting creates a single search head.) For guidance, see Search Head Clustering earlier in this guide.
Apps/Add-ons to pre-install on Splunk indexers (IndexerApps)	—	Comma-separated list of URLs for the Splunk app (or add-on) tarballs (.spl files) to pre-install on indexer(s). For more information about these, see step 4.
Apps/Add-ons to pre-install on Splunk search heads (SearchHeadApps)	—	Comma-separated list of URLs for the Splunk app (or add-on) tarballs (.spl files) to pre-install on search head(s). For more information about these, see step 4.

AWS Quick Start Configuration:

Parameter label (name)	Default	Description
Quick Start S3 Bucket Name (QSS3BucketName)	aws-quickstart	S3 bucket where the Quick Start templates and scripts are installed. Use this parameter to specify the S3 bucket name you've created for your copy of Quick Start assets, if you decide to customize or extend the Quick Start for your own use. The bucket name can include numbers, lowercase letters, uppercase letters, and hyphens, but should not start or end with a hyphen.
Quick Start S3 Key Prefix (QSS3KeyPrefix)	quickstart-splunk-enterprise/	The S3 key name prefix used to simulate a folder for your copy of Quick Start assets, if you decide to customize or extend the Quick Start for your own use. This prefix can include numbers, lowercase letters, uppercase letters, hyphens, and forward slashes.

• Option 2: Parameters for deploying Splunk Enterprise into an existing VPC

View template

AWS Instance and Network Settings:

Parameter label (name)	Default	Description
EC2 instance type for Splunk indexer (IndexerInstanceType)	c5.4xlarge	EC2 instance type to use for Splunk Enterprise indexers. For guidance on selecting an instance type, see Instance Selection earlier in this guide.
EC2 instance type for Splunk search head (SearchHeadInstance Type)	c5.4xlarge	EC2 instance type to use for Splunk Enterprise search heads. For guidance on selecting an instance type, see Instance Selection earlier in this guide.

Parameter label (name)	Default	Description
Key Name (KeyName)	*Requires input*	Public/private key pair, which allows you to connect securely to your instance after it launches. When you created an AWS account, this is the key pair you created in your preferred region.
Permitted CIDR for Splunk web interface (WebClientLocation)	*Requires input*	The CIDR IP range that is permitted to access the Splunk servers' web interface. We recommend that you set this value to a trusted IP range.
Permitted CIDR for Splunk HTTP event collector input (HECClientLocation)	*Requires input*	The CIDR IP range that is permitted to access Splunk HTTP Event Collector (HEC). We recommend that you set this value to a trusted IP range.
Permitted CIDR for SSH (SSHClientLocation)	*Requires input*	The CIDR IP range that is permitted to access Splunk Enterprise instances via SSH. We recommend that you set this value to a trusted IP range. For example, you might want to grant only your corporate network access to the Splunk deployment.
VPC ID (VPCID)	*Requires input*	ID of your existing VPC (e.g., vpc-0343606e).
VPC CIDR (VPCIDR)	*Requires input*	The CIDR block for your existing VPC (e.g., 10.0.0.0/16).
Public Subnet 1 ID (PublicSubnet1ID)	*Requires input*	ID of the public subnet in Availability Zone 1 in your existing VPC (e.g., subnet-a0246dcd).
Public Subnet 2 ID (PublicSubnet2ID)	*Requires input*	ID of the public subnet in Availability Zone 2 in your existing VPC (e.g., subnet-b60d4e78).
Public Subnet 3 ID (PublicSubnet3ID)	—	(Optional) ID of the public subnet in Availability Zone 3 in your existing VPC (e.g., subnet-b58c3d67).
Number of Availability Zones (NumberOfAZs)	2	The number of Availability Zones to use in the VPC. This must match your selections in the Availability Zones parameter. You can choose 2 or 3 Availability Zones.

Splunk Settings:

Parameter label (name)	Default	Description
Splunk Admin Password (SplunkAdminPassword)	*Requires input*	The password for Splunk Enterprise. This string must be 6-32 characters long, and may contain letters, numbers, and symbols.
Shared Security Key for Cluster Nodes (SplunkClusterSecret)	*Requires input*	Shared cluster secret for the Splunk search head and indexer clusters. This string must be 8-32 characters long, and may contain letters, numbers, and symbols.
Shared Security Key for Forwarders using	*Requires input*	Security key used for communications between forwarders and the cluster master. This value should also be used by

Parameter label (name)	Default	Description
Indexer Discovery (SplunkIndexerDiscovery Secret)		forwarders to retrieve a list of available peer nodes from the cluster master. This string must be 8-32 characters long, and may contain letters, numbers, and symbols.
Splunk License Bucket (SplunkLicenseBucket)	—	The name of the private S3 bucket that contains your Splunk license key file, from step 2.
Splunk License S3 Bucket Path (SplunkLicensePath)	—	The path to the S3 bucket that contains your Splunk license key file, without a leading forward slash (/), from step 2.
No. of Splunk Indexers (SplunkIndexerCount)	3	The number of Splunk Enterprise instances to launch. You can choose from 3 to 10 instances.
Indexer Disk Size (SplunkIndexerDiskSize)	320	The size of the EBS volume attached to the Splunk Enterprise indexers, in GiB. You can choose a value between 50-16,000. For guidance, see Instance Selection earlier in this guide.
Search Head Disk Size (SplunkSearchHeadDisk Size)	320	The size of the EBS volume attached to the Splunk search head, in GiB. You can choose a value between 320-16,000. For guidance, see Instance Selection earlier in this guide.
Index Cluster Replication Factor (SplunkReplicationFactor)	2	The number of copies of data to store in the Splunk indexer cluster. You can choose 2-4 copies. For guidance, see Indexers earlier in this guide.
Index Cluster Search Factor (SplunkSearchFactor)	2	The number of copies of data to retain for searching in the Splunk head cluster. You can choose 2-4 copies.
Enable Search Head Cluster? (SHCEnabled)	no	Set this parameter to **yes** to deploy a Splunk search head cluster. (The default setting creates a single search head.) For guidance, see Search Head Clustering earlier in this guide.
Apps/Add-ons to pre-install on Splunk indexers (IndexerApps)	—	Comma-separated list of URLs for the Splunk app (or add-on) tarballs (.spl files) to pre-install on indexer(s). For more information about these, see step 4.
Apps/Add-ons to pre-install on Splunk search heads (SearchHeadApps)	—	Comma-separated list of URLs for the Splunk app (or add-on) tarballs (.spl files) to pre-install on search head(s). For more information about these, see step 4.

AWS Quick Start Configuration:

Parameter label (name)	Default	Description
Quick Start S3 Bucket Name (QSS3BucketName)	aws-quickstart	S3 bucket where the Quick Start templates and scripts are installed. Use this parameter to specify the S3 bucket name you've created for your copy of Quick Start assets, if you decide to customize or extend the Quick Start for your own use. The bucket name can include numbers, lowercase letters,

Parameter label (name)	Default	Description
		uppercase letters, and hyphens, but should not start or end with a hyphen.
Quick Start S3 Key Prefix (QSS3KeyPrefix)	quickstart-splunk-enterprise/	The S3 key name prefix used to simulate a folder for your copy of Quick Start assets, if you decide to customize or extend the Quick Start for your own use. This prefix can include numbers, lowercase letters, uppercase letters, hyphens, and forward slashes.

5. On the **Options** page, you can specify tags (key-value pairs) for resources in your stack and set advanced options. When you're done, choose **Next**.
6. On the **Review** page, review and confirm the template settings. Under **Capabilities**, select the check box to acknowledge that the template will create IAM resources.
7. Choose **Create** to deploy the stack.
8. Monitor the status of the stack. When the status is **CREATE_COMPLETE**, the Splunk Enterprise cluster is ready.
9. Use the URLs displayed in the **Outputs** tab for the stack to view the resources that were created.

Step 4. Send Data to the Splunk Indexers

When your Splunk deployment has completed successfully, you can start sending data to your indexers. The indexers are already preconfigured to receive data from your universal forwarders (via port 9997) in addition to Splunk HEC (via port 8088).

- If you are developing and running applications, you can have them send application logs or metrics directly to Splunk HEC. Use the stack outputs **HttpEventCollectorURL** and **HttpEventCollectorToken** to send data to the HEC load balancer URL, using the secure token to authenticate traffic. You can also use these HEC values to configure one of the various purpose-built AWS Lambda blueprints for Splunk to stream data and events from AWS. For details, see the Splunk documentation.
- If you have application logs on your instances that need to be sent to Splunk, our general recommendation is to use our universal forwarder with indexer discovery to send data to your indexer cluster. With indexer discovery, each forwarder queries the cluster master for a list of all available peer nodes in the cluster, including any that are later added to the cluster. Use the stack output **ClusterMasterManagementURL** and the security key input you provided with the **SplunkIndexerDiscoverySecret** parameter to configure forwarders with indexer discovery. For details, see the Splunk Enterprise

documentation. For complete documentation about Splunk universal forwarders, see the Splunk Universal Forwarder documentation.

- To ingest and visualize data from many of the most common AWS services, Splunk has built the Splunk Add-On for AWS and Splunk App for AWS. The add-on is the data ingestion mechanism, while the app has dozens of prebuilt dashboards to help you visualize your entire AWS ecosystem.

We highly recommend encrypting all data input and web traffic, by deploying third-party-signed or self-signed SSL certificates. For traffic going through the load balancer (HTTP event collector input and Splunk Web in case of a clustered search head), replace the HTTP listeners for the load balancer with their HTTPS counterparts using your SSL certificates. See the Elastic Load Balancing documentation for detailed steps. For traffic going directly to EC2 instances running Splunk (data forwarding and Splunk Web in case of a standalone search head), see About securing data from forwarders and About securing Splunk Web in the Splunk Enterprise documentation to easily configure Splunk Enterprise to use your SSL certificates.

Troubleshooting

If you encounter a CREATE_FAILED error when you launch the Quick Start, we recommend that you relaunch the template with **Rollback on failure** set to **No**. (This setting is under **Advanced** in the AWS CloudFormation console, **Options** page.) With this setting, the stack's state will be retained and the instance will be left running, so you can troubleshoot the issue. (Look at the log files `/var/log/cloud-init-output.log` and `/var/log/cloud-init.log`)

> **Important** When you set **Rollback on failure** to **No**, you'll continue to incur AWS charges for this stack. Please make sure to delete the stack when you've finished troubleshooting.

For additional information, see Troubleshooting AWS CloudFormation on the AWS website.

Additional Resources

AWS services

- Amazon EC2
  https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/

- AWS CloudFormation
  https://aws.amazon.com/documentation/cloudformation/
- Amazon VPC
  https://aws.amazon.com/documentation/vpc/
- Amazon EBS
  https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AmazonEBS.html
- Elastic Load Balancing
  https://aws.amazon.com/documentation/elastic-load-balancing/

Splunk Enterprise

- Product documentation
  http://docs.splunk.com/Documentation/Splunk/latest/
- Manual implementation guide
  https://www.splunk.com/pdfs/white-papers/splunk-enterprise-on-aws-deployment-guidelines.pdf
- Splunk on AWS technical brief
  https://www.splunk.com/pdfs/technical-briefs/deploying-splunk-enterprise-on-amazon-web-services-technical-brief.pdf
- Splunk Add-on for Amazon Web Services
  https://splunkbase.splunk.com/app/1876/
- Splunk App for AWS
  https://splunkbase.splunk.com/app/1274/
- Splunk HTTP Event Collector
  http://dev.splunk.com/view/event-collector/SP-CAAAE6M
- AWS Lambda blueprints for HEC
  http://dev.splunk.com/view/event-collector/SP-CAAAE6W

Quick Start reference deployments

- AWS Quick Start home page
  https://aws.amazon.com/quickstart/

GitHub Repository

You can visit our GitHub repository to download the templates and scripts for this Quick Start, to post your feedback, and to share your customizations with others.

Document Revisions

Date	Change	In sections
May 2019	Updated parameters with newer instance types, disk sizes, and replication/search factor settings	Instance Selection Parameter table for new VPC Parameter table for existing VPC
September 2017	Added support for three Availability Zones with zone-aware indexer and search head clustering; auto-configuration of HEC across the indexer tier; preconfigured indexer discovery	Changes in templates and throughout guide
February 2017	Initial publication	—

Notices

This document is provided for informational purposes only. It represents AWS's current product offerings and practices as of the date of issue of this document, which are subject to change without notice. Customers are responsible for making their own independent assessment of the information in this document and any use of AWS's products or services, each of which is provided "as is" without warranty of any kind, whether express or implied. This document does not create any warranties, representations, contractual commitments, conditions or assurances from AWS, its affiliates, suppliers or licensors. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.