Service Definition Document

Service Overview

Snyk is a suite of developer and management tools for security scanning, fixing, monitoring and reporting on vulnerabilities in proprietary code, open source vulnerabilities in included libraries and containers, and infrastructure as code (IAC).

Snyk Code (SAST)

Snyk Code is a developer-first next-generation SAST tool, that enables developers to build software securely during development, and not try to find and fix problems after the code is compiled. Snyk Code works in the IDEs and SCMs developers use daily to build and review software and provides fast, actionable, meaningful results to fix issues in real-time.

Unparalleled accuracy

Generally, SAST tools have been notorious for the number of false positives they return. Snyk Code utilizes a semantic analysis AI engine that learns from millions of open-source commits and is paired with Snyk's Security Intelligence database--this creates a continually growing code security knowledge base, which reduces false positives to near-zero and provides actionable findings that matter.

Real-time

Speed is the critical factor if you want to support rapid, agile development. Real-time speed allows developers to leverage Code from the IDE and during code review in the SCM, rather than a slow and unnecessary extra step at the end of the development process. Snyk Code scans 10-50x faster than other SAST products, enabling developers to use it while they develop, rather than after they develop as a slow and disruptive step in their process.

Actionable

Although quickly and accurately detecting potential security flaws in source code is a complicated task, we believe that it's not enough. Snyk can only shift left and empower developers if it actually helps remediate the issue and learn about prevention. Snyk Code leverages its security knowledge base to provide fix examples from real-world projects that offer insight on how to fix the issue. Additionally, Code offers curated educational content about every vulnerability to help developers grow their knowledge and reduce issues over

time.

Open Source Code

Snyk provides a software composition analysis by computing the dependency of an application and cross-referencing that with our proprietary security vulnerability database allowing you to find, fix, and monitor security vulnerabilities.

It also provides license management such that the licenses of all packages used by an application are detected and can be alerted based on predefined and customizable classification.

Snyk Container

Snyk Container enables developers to test, fix and monitor open source vulnerabilities in their images.

Scanning and analyzing your Linux-based container project for known vulnerabilities is an important step in securing your environment by helping you identify and mitigate security vulnerabilities.

To help secure your container, Snyk scans the base image for its dependencies:

- The operating system (OS) packages installed and managed by the package manager
- Key binaries - layers that were not installed through the package manager

Based on the scan results, Snyk offers remediation advice and guidance for public DockerHub images by indicating the:

- Origins of the vulnerabilities in your OS packages and key binaries
- Base image upgrade details or a recommendation to rebuild the image
- Dockerfile layer in which the affected package was introduced
- Fixed-in version of the operating system and key binary packages

Developers can also leverage container image scanning from the CLI, enabling a shift left in secure Docker image development.

Containers Orchestrators

Container orchestration is another area that has moved from being an operations concern to

a developer one and is an area where there is the potential to introduce security issues.

Kubernetes is quite complex in its level of configurability of clusters and some of the default values are not the most secure.

Snyk detects Kubernetes configuration files and scans for common configuration issues. These issues are displayed alongside the offending line/s in the configuration file allowing you to quickly understand and remediate the problems.

Snyk IAC

Snyk Infrastructure as Code (Snyk IaC) helps developers write secure HashiCorp Terraform, AWS CloudFormation, Kubernetes, and Azure Resource Manager (ARM) configurations before touching production. Snyk's developer-first approach meets developers where they work and provides fixes that can be directly merged into code.

Information assurance

The services are aligned with good security practices and are appropriate for the processing and storage of information marked OFFICIAL.

On-boarding

During the discovery phase, customers will work with the sales team to define the scope and plan a priorities integration plan. We generally communicate with customers over in person / video meetings, shared slack channels and email.

This plan will then be executed with the customer being assigned a customer success team who will provide implementation instructions and help, training sessions and best practice advice.

Once the customer is live with the solution the customer success manager will be assigned to be a key contact for support and ongoing support for existing and new features of the platform.

Off-boarding

An off boarding customer can export all reports and audit logs before closing the account.

All customer data will then be removed from the platform and deleted.

A free account will still be available for the customer to use the free features of the platform.

Training

Snyk Customer Success (CS) team will help you to plan your Snyk adoption, and the path to implementing Snyk into your SDLC based on your policies and your solution architecture, as well as keep you updated with training on new features that are being implemented rapidly.

The CS team will help you to train your team and to build collateral to be used to train new team members.

The CS team will support you through a close, immediate and long term dedicated Customer Success Manager (CSM) who will provide guidance and resolution for any queries and will represent your advocate voice with our product team.

Support is also available through our online knowledge centre which contains searchable and categorised how-to guides, videos and technical documentation.

Pricing

We are licensing our products based on the number of contributing developers per year.

There are Four main products, Snyk Code, Snyk Open Source, Snyk container and Snyk IAC - all available individually or as a bundle.

As an overlay, there are four tiers, "Free", "Team", "Business" and "Enterprise".

"Team" contains the basic functionality including license compliance.

"Business" contains everything as in "Team" plus on-premise source code via a broker (Github, Bitbucket, GitLab), Single Sign-On, Teams & Groups, Jira integration, private registries (Artifactory, Nexus) and Service Accounts.

"Enterprise" contains everything as in "Business" plus SLAs, a dedicated account manager and custom legal terms. In addition, "Enterprise" can be deployed as on-prem (air-gapped). The latter has an extra cost-element.

Backup/restore and disaster recovery

Snyk uses IBM Compose for database server management. To keep the production database consistently running Snyk employs a combination of redundancy, hot-standby, snapshots and point-in-time backups.

Redundancy & Hot-standby

Each Compose deployment includes a mirror replica setup as a hot-standby with data synchronized up to the speed of shipping data increments between GCP data centres. The Compose service is set up for auto-failover allowing a seamless takeover of the hot-backup in case the master instance fails. Furthermore, the mirrored replica resides in a separate GCP availability zone, enabling the continuity of Snyk services even in the unlikely event that one of GCP data centres becomes unavailable.

Snapshots

Full database snapshots are taken daily and stored with a minimum retention of 90 days. Each snapshot is propagated to multiple GCP cloud storages in multiple availability zones and in multiple regions. Daily backups provide a last resort recovery in case of massive data corruption or loss.

Service levels

Snyk is deployed in a scalable environment on the Google Cloud Platform which is easily scalable to ensure the quality of service for all customers.

Snyk offers a 99.9% uptime SLA in any calendar month. This excludes any scheduled downtime. Snyk openly reports availability via https://snyk.statuspage.io/#month

As part of our SLA we:

- Guarantee to restore key services within 4 business hours of the incident
- Recover to business as usual within 8 to 24 hours after the incident.

Termination terms

https://snyk.io/policies/terms-of-service/

Reference section 13 for termination.

Consumer responsibilities

https://snyk.io/policies/acceptable-use-policy/

Technical requirements

The technical requirements required for the platform will depend upon the deployment method.

SaaS

The deployed SaaS solution is very light in terms of client bandwidth requirements and is tolerant to a wide range of latencies.

The command-line tools, which run on a user's PC, are designed to work with a wide range of machine specifications.

Brokered SaaS

The same deployment method as the SaaS solution with the added separation where on-premise SCM tools are connected via a broker container image deployed within the customers' SCM environment.

This Broker mediates communication between the on-premise tools and Snyk

https://support.snyk.io/hc/en-us/articles/360015296277-How-the-Broker-works

Trial service

Snyk offers a free Pilot deployment for clients to evaluate the platform within their own development environment.

This Pilot will be guided by our Sales Engineering team to ensure we have defined the solution requirements and help your teams get up and running very quickly.

We will also open communication channels (Slack, Email) to ensure close contact and any queries or feature requests are handled quickly by the dedicated Sales Engineer.

Data extraction/removal

Extraction

All data related to your projects can be exported from the UI or API:

- Vulnerability reports
- Licence reports
- Dependency reports (BOM)

Audit logs for platform usage can be exported from the UI, API or by sending a request to support@snyk.io.

Removal

Upon termination of services all data related to your organisation (Projects, Connections, User Accounts) will be deleted and access to the platform will be removed.

Data storage and processing locations

The SaaS platform is hosted in the United States on the Google Cloud Platform - GCP US East 1 with Privacy shield.

Deployment / Service models

Snyk can be delivered as SaaS or Brokered SaaS model.

We are currently developing a private cloud deployment for customers that require it

Elastic resources

To prevent reduced or limited access to the various Snyk services, server pools automatically expand and shrink to meet request demands using Auto Scaling Groups, and traffic is distributed between the servers using elastic load balancers.

Persistence of storage

Data is persisted in the Snyk platform according to the below specification:

- Postgres Database Daily Snapshots Minimum 90d
- Production Application Log Files Minimum 90d
- Production Audit Logs Minimum 90d

If persisted data is required for longer periods of time

Service provisioning

The entire suite of Snyk's application host pools is provisioned by scripted automated infrastructure launching tools. This enables the building or relocation of a deployed

environment to be as simple as a click of a button while being managed as code for proper change management and audit purposes.

Utilisation monitoring/reporting

The Reports area of the UI offers data and analytics across all of your projects, displaying historical and aggregated data about projects, issues, dependencies, and licenses.

Data is displayed based on the organization in which you are working, and you can filter this data with different parameters depending on the data set you are viewing.

Additionally, if your account is managed with groups, aggregated data for all of your organizations is displayed when you navigate to Reports from the Group level. From this Group level, you can filter to view data for multiple organizations

All reports can be exported and generated with our APIs.

Data centre(s)

United States, Google Cloud Platform - GCP US East 1 with Privacy shield.

Service roadmaps

Available on request