

# SPLK-1003 Training Course

Splunk Enterprise Certified Admin

Structured Learning & Certification Preparation

# Table of Contents

<a href="#">SPLK-1003 Training Course</a>	1
<a href="#">Splunk Enterprise Certified Admin</a>	1
<a href="#">Structured Learning &amp; Certification Preparation</a>	1
<a href="#">Table of Contents</a>	2
<a href="#">Introduction</a>	5
<a href="#">About This Training / Certification</a>	5
<a href="#">What We Offer (AAAdemy)</a>	5
<a href="#">Knowledge Overview</a>	6
<a href="#">Detailed Knowledge Explanation</a>	8
<a href="#">SPLK-1003 Splunk Admin Basics</a>	8
<a href="#">1. Core Components of Splunk</a>	8
<a href="#">1.1 Search Head</a>	9
<a href="#">1.2 Indexer</a>	9
<a href="#">1.3 Forwarders</a>	9
<a href="#">1.4 Cluster Manager</a>	9
<a href="#">1.5 Deployment Server</a>	9
<a href="#">2. Splunk Data Pipeline</a>	9
<a href="#">3. Basic Administration Tasks</a>	10
<a href="#">4. Configuration Precedence and Directory Priority</a>	10
<a href="#">5. Splunk Admin Basics Practice Question</a>	10
<a href="#">SPLK-1003 Getting Data In</a>	12
<a href="#">1. Input Methods</a>	12
<a href="#">2. Metadata Assignment</a>	12
<a href="#">3. Getting Data In Practice Question</a>	12
<a href="#">SPLK-1003 Monitor Inputs</a>	14
<a href="#">1. Monitorable Sources</a>	14
<a href="#">2. Performance Tuning</a>	14
<a href="#">3. CRC Mechanism and Duplicate Prevention</a>	14
<a href="#">4. Monitor Inputs Practice Question</a>	14
<a href="#">SPLK-1003 Network and Scripted Inputs</a>	16
<a href="#">1. Network Inputs</a>	16
<a href="#">2. Scripted Inputs</a>	16
<a href="#">3. Network and Scripted Inputs Practice Question</a>	16
<a href="#">SPLK-1003 Agentless Inputs</a>	18
<a href="#">1. Input Methods</a>	18
<a href="#">2. Key Considerations</a>	18
<a href="#">3. Agentless Inputs Practice Question</a>	18
<a href="#">SPLK-1003 Configuring Forwarders</a>	20
<a href="#">1. Types of Forwarders</a>	20
<a href="#">2. Configuration Tasks</a>	20
<a href="#">3. Configuring Forwarders Practice Question</a>	21

<u>SPLK-1003 Forwarder Management</u>	<u>22</u>
<u>1. Centralized Management</u>	<u>22</u>
<u>2. Intermediate Forwarders</u>	<u>22</u>
<u>3. Monitoring and Optimization</u>	<u>22</u>
<u>4. Forwarder Management Practice Question</u>	<u>23</u>
<u>SPLK-1003 Distributed Search</u>	<u>24</u>
<u>1. Distributed Environment Components</u>	<u>24</u>
<u>2. Configuration and Optimization</u>	<u>24</u>
<u>3. Distributed Search Practice Question</u>	<u>25</u>
<u>SPLK-1003 Splunk Indexes</u>	<u>26</u>
<u>1. Index Types</u>	<u>26</u>
<u>2. Bucket Lifecycle</u>	<u>26</u>
<u>3. Managing Indexes</u>	<u>27</u>
<u>4. Splunk Indexes Practice Question</u>	<u>27</u>
<u>SPLK-1003 License Management</u>	<u>28</u>
<u>1. License Types and Roles</u>	<u>28</u>
<u>2. License Pools and Stacks</u>	<u>29</u>
<u>3. Violations and Enforcement</u>	<u>29</u>
<u>4. License Management Practice Question</u>	<u>29</u>
<u>SPLK-1003 Parsing Phase and Data</u>	<u>30</u>
<u>1. Parsing Overview</u>	<u>31</u>
<u>2. Index-Time vs. Search-Time Extraction</u>	<u>31</u>
<u>3. Parsing Phase and Data Practice Question</u>	<u>31</u>
<u>SPLK-1003 Manipulating Raw Data</u>	<u>32</u>
<u>1. Configuration Basics</u>	<u>33</u>
<u>2. Advanced Techniques</u>	<u>33</u>
<u>3. TRANSFORMS vs. REPORT</u>	<u>33</u>
<u>4. Manipulating Raw Data Practice Question</u>	<u>33</u>
<u>SPLK-1003 Fine Tuning Inputs</u>	<u>35</u>
<u>1. Optimization Techniques</u>	<u>35</u>
<u>2. Reliability vs. Performance</u>	<u>35</u>
<u>3. Fine Tuning Inputs Practice Question</u>	<u>35</u>
<u>SPLK-1003 Getting Data In – Staging</u>	<u>36</u>
<u>1. Staging Environment Overview</u>	<u>37</u>
<u>2. Validation Checks</u>	<u>37</u>
<u>3. Getting Data In – Staging Practice Question</u>	<u>37</u>
<u>SPLK-1003 Splunk User Management</u>	<u>38</u>
<u>1. User Roles and Capabilities</u>	<u>38</u>
<u>2. Custom Roles and Access Control</u>	<u>39</u>
<u>3. Splunk User Management Practice Question</u>	<u>39</u>
<u>SPLK-1003 Splunk Authentication Management</u>	<u>40</u>
<u>1. Authentication Methods</u>	<u>40</u>
<u>2. Configuration and Troubleshooting</u>	<u>41</u>

<a href="#">3. Splunk Authentication Management Practice Question</a>	41
<a href="#">SPLK-1003 Splunk Configuration Files</a>	42
<a href="#">1. Critical Configuration Files</a>	42
<a href="#">2. Maintenance and Debugging</a>	42
<a href="#">3. Splunk Configuration Files Practice Question</a>	43
<a href="#">Learning Path &amp; Study Advice</a>	44
<a href="#">Who This PDF Is For</a>	45
<a href="#">Call To Action</a>	45

## Introduction

The SPLK-1003 Splunk Enterprise Certified Admin certification is intended to reflect the knowledge and practical understanding required to administer a Splunk Enterprise environment. It represents the ability to work with core administrative components such as platform configuration, data onboarding, access management, indexing, and distributed search. In a modern IT context, this certification is relevant for professionals who support log management, operational visibility, security monitoring, and data-driven troubleshooting across enterprise systems.

## About This Training / Certification

This certification is generally positioned at an intermediate administrative level. It is designed for learners who already understand basic Splunk search and platform concepts and are moving into day-to-day administration responsibilities. The knowledge scope emphasizes how Splunk Enterprise is configured, maintained, and managed in operational environments. It fits naturally into a broader learning path that starts with foundational platform use, then advances into administration, and later expands into architecture, scaling, deployment design, and specialized operational practices.

## What We Offer (AAAdemy)

AAAdemy provides structured training resources designed to support certification preparation and skill development across a wide range of IT domains. Our learning materials are built around clear knowledge structures, practical study guidance, and exam-oriented practice to help learners progress with confidence.

We offer well-organized knowledge explanations that break down complex topics into clear, understandable sections aligned with official exam objectives and real-world skill requirements. Each topic is designed to support both conceptual understanding and practical application.

Our study plans and learning guidance help learners follow a logical progression, focusing on key concepts, common pitfalls, and effective preparation strategies. This approach enables learners to study efficiently while maintaining a clear view of their learning goals.

To reinforce understanding, AAAdemy also provides practice questions and exam-focused insights that reflect typical certification scenarios. These resources are intended to help learners evaluate their readiness and strengthen their confidence before taking an exam.

All content is designed for flexible, self-paced learning, allowing individuals to study independently or alongside their existing professional or academic commitments.

# Knowledge Overview

## Area: Splunk Administration Fundamentals

Candidates are expected to understand the core administrative foundations of Splunk Enterprise. This includes the basic structure of the platform, the role of the administrator, and the relationship between search, indexing, configuration, and data ingestion. A strong conceptual grasp of administrative workflows is important because many advanced tasks build directly on these fundamentals.

## Area: License Management

This area covers the purpose and operational importance of licensing in Splunk Enterprise. Candidates should understand how license models affect data indexing and platform usage, how license pools and stacks are organized, and why license compliance matters in ongoing administration. The focus is on understanding how licensing supports platform governance and continuity rather than treating it as an isolated configuration task.

## Area: Splunk Configuration Files

Candidates should understand how Splunk uses configuration files to control system behavior. This includes the layered configuration model, the relationship between default and local settings, and the importance of configuration precedence across apps and system contexts. Administrators are expected to recognize how settings are applied, how configuration changes affect services, and why disciplined configuration management is central to a stable deployment.

## Area: Splunk Indexes

This domain focuses on how data is logically organized and stored in Splunk Enterprise. Candidates should understand the purpose of indexes, how indexes support search and retention strategy, and how index design affects performance, storage planning, and data access. The emphasis is on understanding the role of indexes in a broader data lifecycle rather than on isolated syntax or memorization.

## Area: Splunk User Management

Candidates are expected to understand how users are created and managed within Splunk Enterprise. This includes concepts related to roles, permissions, inherited capabilities, and access boundaries. The goal is to understand how administrative control supports security, usability, and appropriate separation of responsibilities across teams.

## Area: Splunk Authentication Management

This area covers how Splunk Enterprise verifies identity and controls access. Candidates should understand native authentication concepts as well as the general purpose of integrating enterprise authentication methods. The focus is on how authentication mechanisms support secure administration, centralized identity practices, and consistent access control across environments.

## Area: Getting Data In

This domain addresses the general process of onboarding data into Splunk Enterprise. Candidates should understand the types of data sources that can be collected, the importance of source planning, and the relationship between ingestion design and downstream search quality. This includes conceptual awareness of how data inputs influence parsing, indexing, and operational value.

## Area: Distributed Search

Candidates should understand the purpose and structure of distributed search in Splunk environments. This includes how search heads and indexers interact, how knowledge and search requests are coordinated, and why distributed search is important for scale, performance, and multi-system visibility. The emphasis is on architectural understanding and operational awareness.

## Area: Getting Data In - Staging

This area focuses on the planning and preparation that often precede production data onboarding. Candidates should understand why staging is useful for validating source behavior, data quality, parsing expectations, and configuration outcomes before broad rollout. This reflects an administrative mindset that prioritizes control, predictability, and reduced operational risk.

## Area: Configuring Forwarders

Candidates are expected to understand the role of forwarders in collecting and sending data to Splunk Enterprise. This includes how forwarders fit into the ingestion architecture, what types of forwarding behavior exist, and why proper configuration is important for reliability, scale, and maintainability. Forwarders should be understood not only as data shippers but as key components in enterprise collection strategy.

## Area: Forwarder Management

This domain covers the operational management of deployed forwarders. Candidates should understand how administrators maintain consistency, monitor forwarder behavior, and apply changes across multiple systems. The focus is on lifecycle management, operational control, and the importance of standardization in larger environments.

## Area: Monitor Inputs

Candidates should understand how file and directory monitoring supports continuous data collection. This includes awareness of how monitored sources are tracked, how changes are detected, and why input design matters for completeness and efficiency. The emphasis is on ensuring that data is collected in a dependable and predictable manner.

## Area: Network and Scripted Inputs

This area covers data ingestion methods that rely on network services or scripted collection. Candidates should understand the purpose of these input types, the kinds of use cases they support, and the operational considerations involved in reliability, formatting, and execution context. This reflects the need to manage data sources that are more dynamic than simple file monitoring.

#### Area: Agentless Inputs

Candidates should understand the concept of collecting data without deploying an agent directly on the source system. The focus is on why agentless methods may be used, what trade-offs they introduce, and how they fit into broader ingestion design decisions. This area highlights the importance of choosing collection methods that align with technical, security, and operational constraints.

#### Area: Fine Tuning Inputs

This domain focuses on improving data quality and collection behavior through input refinement. Candidates should understand why tuning matters for event accuracy, source consistency, noise reduction, and downstream usability. The goal is to recognize that effective administration involves not only enabling inputs, but also optimizing them so the resulting data is meaningful and efficient to process.

#### Area: Parsing Phase and Data

Candidates are expected to understand the role of parsing in the Splunk data pipeline. This includes how raw incoming data is interpreted, segmented into events, and prepared for indexing. A conceptual understanding of the parsing phase is important because it directly affects search quality, field extraction behavior, and the overall usability of indexed data.

#### Area: Manipulating Raw Data

This area addresses how raw incoming data may be adjusted or transformed before it becomes fully searchable content. Candidates should understand the purpose of such manipulation, the administrative reasons for applying it, and the care required to preserve data integrity and search usefulness. The emphasis is on controlled handling of incoming data to improve normalization, readability, and operational value.

## Detailed Knowledge Explanation

### **SPLK-1003 Splunk Admin Basics**

Splunk Enterprise utilizes a modular and distributed architecture that allows organizations to scale data processing capabilities from a single instance to a global cluster. By decoupling the stages of the data lifecycle—collection, parsing, indexing, and search—Splunk ensures that system resources are allocated precisely to maintain high-performance ingestion alongside complex analytical workloads. This structural flexibility is fundamental to providing real-time operational visibility across diverse IT environments, where components must be managed through both the Splunk Web interface and the Command Line Interface (CLI).

#### **1. Core Components of Splunk**

The strategic separation of search and storage functions is critical for system stability in high-volume environments. By isolating user-facing query tasks from heavy data storage operations, Splunk prevents search

activities from impacting the continuous ingestion of raw data. This distributed model relies on specialized components that work in concert to move data through its lifecycle while maintaining data integrity and system responsiveness.

### 1.1 Search Head

The Search Head serves as the primary user interface for the Splunk environment, acting as a query manager that does not store raw data. Its primary responsibility is distributing Search Processing Language (SPL) queries to the indexing tier, aggregating the returned results, and presenting them through dashboards, reports, and alerts. In large-scale deployments, multiple Search Heads can form a Search Head Cluster to provide redundancy and handle high query loads, requiring synchronization of knowledge objects such as saved searches and lookups.

### 1.2 Indexer

The Indexer is the backbone of the Splunk platform, responsible for transforming raw data into searchable events and persisting them to disk. During the ingestion process, the Indexer tokenizes data and assigns essential metadata such as timestamps and sourcetypes. Data is organized into logical storage units known as buckets, which transition through a lifecycle of Hot, Warm, Cold, and Frozen states. This bucket system allows administrators to optimize storage costs by moving older data to more economical storage tiers while maintaining searchability.

### 1.3 Forwarders

Forwarders are lightweight agents installed on data sources to collect and transmit data to the indexing tier. The Universal Forwarder (UF) is a resource-efficient option designed for high-performance transmission with minimal overhead, making it ideal for massive deployments on production servers. Conversely, the Heavy Forwarder (HF) is a full Splunk instance capable of parsing, filtering, and masking data before transmission. This allows for advanced preprocessing at the source, though it requires significantly more system resources than a Universal Forwarder.

### 1.4 Cluster Manager

The Cluster Manager acts as the central orchestrator for Indexer Clusters, maintaining high availability and data redundancy across the distributed environment. It coordinates data replication and ensures that the Replication Factor (RF) and Search Factor (SF) are met across the cluster nodes. By monitoring the state of all peers, the Cluster Manager facilitates automatic failover and ensures that data remains searchable and resilient even during hardware failures.

### 1.5 Deployment Server

The Deployment Server provides centralized configuration management for a fleet of Splunk instances, particularly Universal Forwarders. Administrators use it to package configurations into deployment apps and distribute them to specific groups known as server classes. This model simplifies the management of thousands of forwarders by providing a single point of visibility to monitor "phone home" status and ensure that agents are running the most current **inputs.conf** and **outputs.conf** settings.

## 2. Splunk Data Pipeline

Splunk processes data through a structured four-stage pipeline consisting of Input, Parsing, Indexing, and Search. In the Input stage, data is collected from sources like files or network ports. The Parsing stage acts as the "brain" where the system breaks the raw stream into individual events, identifies timestamps, and assigns metadata. During the Indexing stage, parsed events are converted into a searchable format and written to disk. Finally, the Search stage enables users to utilize SPL to query and visualize the indexed data for real-time analysis.

### 3. Basic Administration Tasks

The lifecycle of Splunk management involves service control and continuous health monitoring using the CLI and the Monitoring Console. Administrators use commands such as **splunk start**, **splunk stop**, and **splunk restart** to manage the service lifecycle and apply configuration changes. Health monitoring is conducted through the Monitoring Console, which tracks indexing performance and search latency, and through internal logs like **splunkd.log** and **metrics.log** located in **\$SPLUNK\_HOME/var/log/splunk/**.

### 4. Configuration Precedence and Directory Priority

Splunk utilizes a sophisticated directory hierarchy where local settings override defaults and system settings generally override app-specific layers. The precedence order follows a path from **etc/system/local/** down through **etc/apps/local/**, **etc/apps/default/**, and finally **etc/system/default/**. To troubleshoot effective configurations, administrators use the **splunk btool** utility with the **--debug** flag, which provides a merged view of all active settings and identifies the exact file path source for every parameter, ensuring system integrity is maintained across the environment.

Having established the architectural foundation, the following section details the specific methodologies for ingesting data.

### 5. Splunk Admin Basics Practice Question

Q1: Which Splunk component is responsible for executing search queries, generating dashboards, and creating alerts, but does not store data?

- A. Indexer
- B. Forwarder
- C. Search Head
- D. Cluster Manager

Q2: What is the main role of a Deployment Server in a Splunk environment?

- A. It ingests raw data from various sources and writes it to indexes
- B. It pushes configuration updates to forwarders and manages deployment apps
- C. It monitors the cluster health and manages search head replication
- D. It acts as a centralized license manager for forwarders

Q3: Which type of Splunk Forwarder is optimized for lightweight data collection and minimal system resource usage?

- A. Heavy Forwarder
- B. Intermediate Forwarder

- C. Edge Forwarder
- D. Universal Forwarder

Q4: Which Splunk component is responsible for storing data, indexing it, and making it searchable?

- A. Search Head
- B. Indexer
- C. Forwarder
- D. Deployment Server

Q5: During installation, which of the following is the default port for accessing Splunk Web?

- A. 443
- B. 9997
- C. 8089
- D. 8000

Q6: In which stage of the Splunk data pipeline are events assigned timestamps and broken into individual entries?

- A. Input Stage
- B. Indexing Stage
- C. Search Stage
- D. Parsing Stage

Q7: Which CLI command is used to validate and debug configuration file settings in Splunk?

- A. splunk btool
- B. splunk diag
- C. splunk show config
- D. splunk start

Q8: In which stage of the Splunk data pipeline is raw data written to index buckets for storage and searchability?

- A. Indexing Stage
- B. Input Stage
- C. Search Stage
- D. Parsing Stage

Q9: Which log file should an administrator check for general Splunk errors and service issues?

- A. metrics.log
- B. license.log
- C. splunkd.log
- D. audit.log

Q10: What is the purpose of using a Heavy Forwarder instead of a Universal Forwarder in a Splunk deployment?

- A. To perform data filtering, masking, and transformation before indexing
- B. To manage clustering and replication factors
- C. To simplify log collection from mobile devices
- D. To reduce memory usage on Indexers

## SPLK-1003 Getting Data In

Modern IT ecosystems generate a vast array of data formats, necessitating diverse ingestion strategies to maintain comprehensive visibility. Splunk accommodates this by providing mechanisms for monitoring persistent logs, receiving real-time network streams, and fetching data through scripts or APIs. Selecting the appropriate ingestion method is a strategic decision that balances system performance with the requirements for data accuracy and low-latency visibility across the enterprise.

### 1. Input Methods

Splunk provides several primary mechanisms to ingest data according to its delivery requirements. File and directory monitoring is the standard for local logs, where Splunk tracks modifications and ingests new data in real time. Network inputs allow Splunk to listen on TCP or UDP ports to receive streams like Syslog. Scripted inputs execute custom code at defined intervals to fetch metrics or API data, while the HTTP Event Collector (HEC) offers a high-performance, token-based API for cloud-native applications to push JSON data directly to Splunk using the **batchMode** setting to optimize throughput.

### 2. Metadata Assignment

The strategic assignment of Host, Source, and Sourcetype metadata is essential for the efficient categorization and retrieval of data. The Host field identifies the originating system, while the Source specifies the file path, port, or script that generated the data. The Sourcetype is particularly critical as it determines how Splunk parses the raw data into searchable events. Proper classification ensures that data is easily filtered during searches and that specific parsing rules in **props.conf** are applied consistently across the platform.

Beyond general ingestion, specific monitoring techniques are required for persistent local data sources.

### 3. Getting Data In Practice Question

Q1: What is the purpose of the `inputs.conf` file in Splunk?

- A. To define and configure data input sources
- B. To store indexed data
- C. To map LDAP groups to roles
- D. To extract fields at search time

Q2: What does the `sourcetype` attribute control in Splunk data ingestion?

- A. The location where data is stored
- B. The timestamp assignment of events
- C. How the data is parsed and fields are extracted
- D. The index used for storing the data

Q3: Which field identifies the type of data and how Splunk should parse it?

- A. host
- B. source

- C. sourcetype
- D. index

Q4: Which Splunk input method is used for collecting real-time data over UDP?

- A. Scripted Input
- B. File Monitoring
- C. HTTP Event Collector
- D. Network Input

Q5: What is the best method to send JSON-formatted log data from an application to Splunk?

- A. Monitor a local file
- B. HTTP Event Collector (HEC)
- C. Scripted Input
- D. TCP Input

Q6: What is the purpose of setting `recursive = true` in a monitor stanza?

- A. To enable automatic field extraction
- B. To include subdirectories in file monitoring
- C. To index the data in real-time
- D. To exclude symbolic links

Q7: Which configuration enables Splunk to listen for events on a specific TCP port?

- A. `[monitor:///var/log/messages]`
- B. `[udp://10514]`
- C. `[script://fetch_data.py]`
- D. `[tcp://10514]`

Q8: What field in metadata determines the hostname of the originating system?

- A. host
- B. source
- C. sourcetype
- D. index

Q9: Which SPL query shows data ingestion throughput by host?

- A. `index=_internal source=*license_usage.log* | stats count by host`
- B. `index=_audit action=login | stats count by host`
- C. `index=_internal source=*metrics.log group=per_host_thruput | stats sum(kbps) by host`
- D. `index=main sourcetype=server_logs | stats avg(kbps) by host`

Q10: What is the purpose of the HTTP Event Collector (HEC) in Splunk?

- A. To configure search time fields
- B. To assign user roles to incoming events
- C. To receive event data over HTTP
- D. To forward data from Indexers to Search Heads

## SPLK-1003 Monitor Inputs

Monitor inputs are the primary method for continuous ingestion of local log data, relying on the Splunk service user's ability to track and read files. Successful implementation requires a deep understanding of file system behaviors and the underlying service permissions. In Linux environments, the Splunk service user must have both read and traverse permissions (r and x) across the target directories, otherwise, the monitor input will show as enabled in the UI but fail to ingest data.

### 1. Monitorable Sources

Splunk monitors individual files or entire directory trees, supporting recursive monitoring to capture data from nested folders. To manage the scope of ingestion and prevent the collection of low-value data, administrators utilize **whitelist** and **blacklist** filters in **inputs.conf**. These filters use regular expressions to include or exclude files based on naming patterns, such as a whitelist for **.log** files combined with a blacklist for files containing the word **debug**, effectively narrowing the ingestion scope.

### 2. Performance Tuning

To optimize resource usage in high-volume environments, administrators can adjust polling intervals to control how frequently Splunk checks for new data in monitored files. Inefficient monitoring of thousands of files can lead to CPU and I/O bottlenecks, so it is best practice to archive or rotate logs that are no longer active and use the **ignoreOlderThan** setting. Limiting the number of open file handles ensures that the system maintains high ingestion throughput without compromising the stability of the host operating system.

### 3. CRC Mechanism and Duplicate Prevention

Splunk utilizes a Cyclic Redundancy Check (CRC) mechanism to prevent duplicate data ingestion by hashing the first 256 bytes of a file. If a file is renamed but the content remains the same, Splunk will skip it to avoid redundancy. To handle rotated logs or files with identical headers, administrators can use the **initCrcLength** setting to expand the hash range or apply **crcSalt =** to include the file path in the uniqueness calculation. These settings are vital for ensuring that moved or renamed files are correctly identified and not ignored as false duplicates.

In addition to local files, Splunk must often act as a receiver for network-based and programmatically generated data.

### 4. Monitor Inputs Practice Question

Q1: What is the purpose of the `[monitor://]` stanza in `inputs.conf`?

- A. To monitor Splunk internal logs
- B. To configure indexer clustering

- C. To specify files or directories to ingest
- D. To configure search time field extractions

Q2: In `inputs.conf`, what is the purpose of the `whitelist` setting?

- A. To group inputs for deployment
- B. To include only files matching a pattern for monitoring
- C. To exclude files from indexing
- D. To define indexer destinations

Q3: Which index contains internal logs and throughput metrics for monitoring input performance?

- A. `_metrics`
- B. `_audit`
- C. `_introspection`
- D. `_internal`

Q4: What happens if `crcSalt` is not set for files with similar headers?

- A. Splunk assigns the wrong sourcetype
- B. Splunk fails to start
- C. Splunk indexes the file twice
- D. Splunk may skip ingestion assuming the file is a duplicate

Q5: You want to ingest `.log` files but exclude `debug.log`. Which blacklist is appropriate?

- A. `blacklist = *`
- B. `blacklist = .debug.`
- C. `blacklist = debug.log$`
- D. `blacklist = .log$`

Q6: What is the effect of setting `followTail = true` in a monitor input?

- A. Ingest only new events from the end of the file
- B. Prevent log rotation
- C. Force recursive monitoring
- D. Schedule inputs on a timer

Q7: You want to include `.json` files but exclude `debug.json`. Which configuration is correct?

- A. `whitelist = .json$ blacklist = debug.json$`
- B. `whitelist = .json$ blacklist = *`
- C. `whitelist = .json. blacklist = .debug.`
- D. `whitelist = *.log$ blacklist = debug.log$`

Q8: What SPL query best verifies input throughput for monitored files?

- A. `index=_audit sourcetype=syslog`
- B. `index=main | stats count by sourcetype`
- C. `index=_internal source=*metrics.log group=per_host_thruput`
- D. `index=_internal source=*splunkd.log error`

Q9: Which setting enables recursive monitoring of all subdirectories under a given path?

- A. `readDepth = max`

- B. recursive = true
- C. directoryDepth = 10
- D. multiFile = true

Q10: What does setting `interval = 300` in `inputs.conf` do?

- A. Sets a 5-minute delay between polling the input
- B. Disables the input after 300 seconds
- C. Limits indexing rate
- D. Sets maximum event age

## SPLK-1003 Network and Scripted Inputs

Capturing data in real time from network devices and custom external sources requires robust communication protocols that transcend simple file monitoring. Network and scripted inputs enable Splunk to serve as a central hub for telemetry, cloud logs, and system metrics. These methods are especially valuable for securing infrastructure where installing local agents is prohibited, allowing for the ingestion of critical logs via standardized protocols.

### 1. Network Inputs

Network collection typically involves Syslog over UDP or TCP, as well as the HTTP Event Collector (HEC). While UDP is lightweight, TCP is preferred for mission-critical data because it provides guaranteed delivery. For cloud-native environments, HEC offers a scalable, token-based method for pushing data via HTTPS. Utilizing features like **enableAck** (Index Acknowledgment) with HEC ensures that data is not lost, as the client can poll the `/ack` endpoint to confirm that the Indexer has successfully written the events to disk.

### 2. Scripted Inputs

Scripted inputs provide a flexible way to fetch data from sources lacking traditional logging mechanisms, such as web APIs. These scripts reside in `$$SPLUNK_HOME/bin/scripts/` and are executed at defined intervals. Best practices for scripted inputs emphasize script efficiency and error logging to `_internal`. Administrators should test scripts manually to ensure they are executable and that their output is formatted correctly, typically as one event per line in JSON, to facilitate downstream parsing and field extraction.

For environments where installing local agents is restricted, agentless ingestion provides a critical alternative.

### 3. Network and Scripted Inputs Practice Question

Q1: What is the function of the interval setting in a scripted input configuration?

- A. Determines max data size
- B. Specifies log rotation policy

- C. Sets how often the script runs
- D. Controls how many tokens are used

Q2: Why should you use a dedicated Syslog server like rsyslog before forwarding data to Splunk?

- A. To avoid using Splunk entirely
- B. To apply HEC token authorization
- C. To improve reliability and centralize collection
- D. To filter JSON logs only

Q3: Which SPL query is useful for monitoring HEC ingestion failures?

- A. `index=_internal sourcetype=splunkd component=HttpEventCollector`
- B. `index=_audit action=login`
- C. `index=hec_data sourcetype=json`
- D. `index=_introspection component=Indexer`

Q4: What is a recommended security measure for HEC data collection?

- A. Send tokens over UDP
- B. Use base64-encoded credentials
- C. Enable SSL/TLS for HEC endpoints
- D. Open port 22 for all data sources

Q5: What is the default port used by the HTTP Event Collector (HEC) in Splunk?

- A. 9997
- B. 514
- C. 8088
- D. 10514

Q6: Which index contains logs from the HTTP Event Collector component for troubleshooting?

- A. `_audit`
- B. `_metrics`
- C. `_hec`
- D. `_internal`

Q7: Which of the following is a valid reason a scripted input may fail to run in Splunk?

- A. Token expiration
- B. Script file missing execution permission
- C. Syslog port misconfiguration
- D. Incorrect sourcetype in transforms.conf

Q8: In which directory should a custom scripted input file be placed?

- A. `$SPLUNK_HOME/bin/scripts/`
- B. `$SPLUNK_HOME/etc/apps/scripts/`
- C. `$SPLUNK_HOME/var/lib/scripts/`
- D. `$SPLUNK_HOME/scripts/`

Q9: What is a best practice when configuring Syslog inputs for critical log data?

- A. Use UDP to reduce latency

- B. Use TCP for reliability
- C. Always listen on port 80
- D. Use a scripted input instead

Q10: Which stanza in inputs.conf is used to configure a scripted input?

- A. [scripted\_input]
- B. [http\_input]
- C. [exec\_input]
- D. [script://]

## SPLK-1003 Agentless Inputs

Agentless ingestion is a strategic necessity in environments with strict security policies or hardware limitations that prevent the installation of a Splunk forwarder. By utilizing existing system protocols and APIs, Splunk can pull data from remote targets without a persistent local footprint. This approach relies on a clear conceptual distinction between push models, where the source initiates data transfer, and pull models, where Splunk initiates the request.

### 1. Input Methods

The primary agentless methods include Windows Management Instrumentation (WMI) and REST API-based ingestion. WMI is a pull model used to collect Windows event logs and performance metrics. Administrators must configure WMI inputs via the Splunk Web UI or REST endpoints and should never use **[script://]** stanzas in **inputs.conf** for WMI, as this is a common configuration error. REST API ingestion, typically implemented via scripted inputs, allows Splunk to poll external services at scheduled intervals, contrasting with the push model of HEC where external systems initiate the transmission.

### 2. Key Considerations

Implementing agentless inputs requires careful attention to security and scalability. Because these methods involve remote communication, securing connections with SSL/TLS or Kerberos is mandatory to protect sensitive credentials and data in transit. Scalability is also a challenge, as frequent remote queries can place significant load on both the network and the target systems. Optimizing WMI queries with **WHERE** clauses and implementing batching in REST scripts are essential techniques to maintain system performance at scale.

While agentless methods are versatile, standardizing agent-based ingestion remains the primary task for distributed log collection.

### 3. Agentless Inputs Practice Question

Q1: Which Splunk input method enables collecting Windows logs without installing a forwarder?

- A. TCP Input

- B. HTTP Event Collector
- C. Modular Input
- D. Windows Management Instrumentation (WMI)

Q2: Which protocol should be used to securely transmit data when using REST API agentless inputs?

- A. HTTPS
- B. HTTP
- C. UDP
- D. FTP

Q3: Which of the following is a common challenge when using WMI for agentless inputs?

- A. Incomplete field extractions
- B. High resource usage on Windows hosts
- C. Inability to schedule queries
- D. Log rotation issues

Q4: What is the default port used for WMI communication (RPC endpoint mapper)?

- A. 9997
- B. 514
- C. 135
- D. 8089

Q5: You need to periodically collect weather data from a public API without a forwarder. Which input type should you configure?

- A. WMI Input
- B. Scripted Input
- C. File Monitor
- D. TCP Input

Q6: Which configuration file defines the interval, sourcetype, and index for a scripted agentless input?

- A. transforms.conf
- B. outputs.conf
- C. inputs.conf
- D. props.conf

Q7: In REST API agentless ingestion, what should be done to reduce the impact of API rate limits?

- A. Use btool to debug the input
- B. Lower the ingestion interval to 1 second
- C. Implement batching and retry logic in the script
- D. Switch to TCP

Q8: How should API credentials be securely managed when using scripted agentless inputs?

- A. Disable SSL
- B. Use environment variables or encrypted credential storage
- C. Store them in a plain text file
- D. Hardcode them in the Python script

Q9: Which of the following commands is used to troubleshoot script execution errors for agentless inputs?

- A. `splunk show deploy-clients`
- B. `python ./bin/scripts/my_script.py`
- C. `splunk list input-status`
- D. `splunk restart`

Q10: What is the main advantage of using WMI over installing Universal Forwarders?

- A. WMI enables lightweight, agentless data collection
- B. WMI supports real-time indexing
- C. WMI is faster for high-volume logs
- D. WMI does not require authentication

## SPLK-1003 Configuring Forwarders

Forwarders are the primary conduits for data in a distributed Splunk architecture, providing reliable delivery of events from the edge to the indexing tier. Properly configured forwarders ensure that data is encrypted, load-balanced, and attributed with correct metadata before transmission. Their role is critical in maintaining the throughput of the data pipeline, especially in environments requiring high-assurance delivery through features like Indexer Acknowledgment and mutual SSL authentication.

### 1. Types of Forwarders

The choice between a Universal Forwarder and a Heavy Forwarder depends on the specific requirements of the data source. The Universal Forwarder is the standard for most deployments, offering a small resource footprint and high efficiency by forwarding raw data without parsing. The Heavy Forwarder is a full Splunk instance reserved for advanced use cases where data must be filtered, parsed, or masked before leaving the source network. While the Heavy Forwarder provides greater control, its increased resource requirements make the Universal Forwarder the preferred choice for massive log collection.

### 2. Configuration Tasks

Configuring forwarders centers on defining data sources in **inputs.conf** and destinations in **outputs.conf**. A key feature is **autoLB**, which enables automatic load balancing across multiple indexers to prevent bottlenecks. For high-security environments, administrators should implement Mutual SSL (Two-Way SSL), where the server also verifies the client's certificate. Furthermore, enabling **useACK = true** in **outputs.conf** ensures that the forwarder maintains events in its queue until the indexer confirms successful receipt and replication.

Managing a fleet of forwarders at scale requires centralized governance and observability.

### 3. Configuring Forwarders Practice Question

Q1: What is the primary function of a Universal Forwarder in a Splunk deployment?

- A. It collects and forwards data without indexing
- B. It stores data in frozen buckets
- C. It parses and indexes data before forwarding
- D. It runs the Splunk Web UI for administration

Q2: What configuration file defines the destination indexers for a Splunk forwarder?

- A. inputs.conf
- B. deploymentclient.conf
- C. outputs.conf
- D. server.conf

Q3: What is the primary role of a Heavy Forwarder compared to a Universal Forwarder?

- A. It provides a UI interface for dashboards
- B. It stores frozen buckets for long-term retention
- C. It has lower resource usage and avoids parsing
- D. It parses, filters, and can locally index data

Q4: Which configuration enables a Universal Forwarder to receive deployment instructions from a Deployment Server?

- A. outputs.conf
- B. deploymentclient.conf
- C. inputs.conf
- D. transforms.conf

Q5: Which output configuration enables data to be sent to multiple indexers with automatic load balancing?

- A. index = cluster\_main
- B. useACK = true
- C. autoLB = true
- D. sslPassword = encrypted

Q6: Which command lists the current forward-server connection status?

- A. splunk list inputs
- B. splunk list forward-server
- C. splunk show deploy-clients
- D. splunk show cluster-status

Q7: Which configuration file should be modified to define what logs a forwarder monitors?

- A. props.conf
- B. outputs.conf
- C. limits.conf
- D. inputs.conf

Q8: When configuring SSL encryption between a forwarder and indexer, which file is commonly used?

- A. transforms.conf

- B. inputs.conf
- C. outputs.conf
- D. indexes.conf

Q9: In the context of forwarders, what is the purpose of a Deployment Server?

- A. To centrally manage configuration deployment to multiple forwarders
- B. To archive frozen data
- C. To search and visualize collected logs
- D. To parse and index incoming data streams

Q10: If logs are being monitored but not appearing in Splunk, which tool helps verify active input settings?

- A. splunk list log-status
- B. splunk check config
- C. splunk cmd btool inputs list --debug
- D. splunk list metadata

## SPLK-1003 Forwarder Management

As forwarder deployments grow into the hundreds or thousands, manual configuration becomes impossible. Centralized management is required to maintain consistency and monitor the health of the collection tier. This governance is handled through the Deployment Server, which provides a unified interface for managing configurations across disparate server groups. Administrators must also account for the license impact, as while forwarders do not require a license to operate, the data they send to indexers is subject to license usage.

### 1. Centralized Management

The Deployment Server uses server classes defined in **serverclass.conf** to group forwarders based on attributes like hostname or IP. Configurations are packaged into deployment apps and pushed to these classes. The Forwarder Management UI in Splunk Web provides real-time visibility into the "last phone home" timestamp for each client and the status of deployed apps. This UI-based management is the standard for deploying configurations to a large fleet and monitoring for missing clients or failed updates.

### 2. Intermediate Forwarders

In complex network topologies, intermediate forwarders serve as relays between edge agents and the indexing tier. They are strategically used to aggregate data from multiple sources, reducing direct connections to the indexers. Intermediate forwarders can also perform advanced data routing and masking using **props.conf** and **transforms.conf**, acting as a gateway that enforces security policies or routes sensitive data to specific regional indexers while dropping low-value logs using the **nullQueue**.

### 3. Monitoring and Optimization

Proactive monitoring of forwarder health is essential for preventing data gaps. Key metrics include throughput, connection status, and deployment health. Administrators use internal logs and the Monitoring Console to identify forwarders with failed connections to indexers or those that have not checked in recently. Optimization techniques, such as enabling data compression with **compressed = true** in **outputs.conf**, can mitigate network bottlenecks, especially in bandwidth-constrained environments.

Distributed ingestion inevitably feeds into the distributed search architecture for final analysis.

#### 4. Forwarder Management Practice Question

Q1: What is the main purpose of the Deployment Server in Splunk?

- A. To store indexed data across multiple indexers
- B. To run SPL queries on Universal Forwarders
- C. To centrally manage configurations across multiple forwarders
- D. To analyze performance metrics of indexers

Q2: In which configuration file do you define a forwarder's connection to a Deployment Server?

- A. inputs.conf
- B. deploymentclient.conf
- C. serverclass.conf
- D. outputs.conf

Q3: Which SPL query helps track the amount of data being sent by forwarders?

- A. index=\_audit | stats count
- B. index=\_internal source=\*metrics.log group=per\_host\_thruput | stats sum(kbps) by host
- C. index=main | stats avg(delay)
- D. index=\_internal | chart avg(cpu\_time)

Q4: What is the function of server classes in the Deployment Server?

- A. To isolate indexers by function
- B. To group forwarders for targeted app deployment
- C. To restrict access to sourcetypes
- D. To configure search head clustering

Q5: What can cause a forwarder to not send data to an indexer?

- A. The sourcetype is incorrect in props.conf
- B. The forwarder's host field is missing
- C. The indexer's SSL certificate is expired
- D. The outputs.conf on the forwarder is misconfigured

Q6: Which command checks the list of connected deployment clients?

- A. splunk show cluster-status
- B. splunk list forward-server
- C. splunk show deploy-clients
- D. splunk list deployment-apps

Q7: What should you configure in outputs.conf to enable load balancing across multiple indexers?

- A. defaultGroup = load\_balance
- B. autoLB = true
- C. sslRootCAPath = auto
- D. indexerBalance = yes

Q8: How can you reduce CPU usage on Universal Forwarders?

- A. Increase maxKBps
- B. Set compressed = false
- C. Reduce the number of monitored inputs
- D. Enable SSL encryption

Q9: What is the purpose of enabling compression in outputs.conf?

- A. It increases read speed from disk
- B. It reduces CPU usage on indexers
- C. It secures connections using TLS
- D. It reduces bandwidth usage during data forwarding

Q10: What is the best way to ensure sensitive configurations on forwarders can be recovered after failure?

- A. Enable acknowledgments in outputs.conf
- B. Use modular inputs
- C. Manually copy inputs.conf to each forwarder
- D. Centralize and back up deployment apps

## SPLK-1003 Distributed Search

Distributed search is the mechanism that enables Splunk to scale horizontally by separating search management from data storage. By executing queries in parallel across multiple indexers, Splunk maintains high performance even as data volumes increase. This architecture requires secure authentication between Search Heads and Indexers (Search Peers) over management port 8089 and proper certificate trust for SSL-enabled environments.

### 1. Distributed Environment Components

The primary components are Search Heads and Indexer Clusters. Data availability in a cluster is governed by the Replication Factor (RF) and the Search Factor (SF). The RF determines how many raw copies of data exist for redundancy, while the SF determines how many are searchable. A Search Head Cluster provides redundancy for the user interface, managed by a Deployer that distributes configuration bundles using the **splunk apply shcluster-bundle** command, though the Deployer itself does not run searches or index data.

### 2. Configuration and Optimization

Connecting Search Heads to indexers is managed through **distsearch.conf** or via the **splunk add search-server** command. For complex queries, administrators can tune **limits.conf** to enable parallel processing, breaking down large searches into concurrent tasks. In multisite clusters, **server.conf** is used to provide geographic redundancy, ensuring that copies of data exist in multiple locations to maintain availability in the event of a site failure.

Central to distributed search is the management of the data itself within the indexing tier.

### 3. Distributed Search Practice Question

Q1: What is the role of a Search Head in a Splunk distributed environment?

- A. It collects data from Universal Forwarders
- B. It replicates indexed data
- C. It manages and executes search queries
- D. It stores raw event data

Q2: What does the replication factor (RF) in an Indexer Cluster determine?

- A. Number of indexers involved in searches
- B. Number of users who can access a report
- C. Number of copies of data stored across the cluster
- D. Number of search heads connected

Q3: Which configuration file is used to connect a Search Head to Indexers?

- A. indexes.conf
- B. distsearch.conf
- C. inputs.conf
- D. server.conf

Q4: In a Search Head Cluster, what is the role of the Deployer?

- A. It pushes configuration bundles to cluster members
- B. It indexes and searches data locally
- C. It forwards data to Indexers
- D. It stores data buckets

Q5: What is the purpose of setting the **max\_searches\_per\_cpu** in limits.conf?

- A. To define the replication factor
- B. To control parallel search execution
- C. To limit log retention
- D. To throttle data ingestion

Q6: What benefit does Search Head clustering provide?

- A. Offers high availability and load balancing for searches
- B. Reduces disk space usage
- C. Enables data forwarding
- D. Increases bucket replication

Q7: Which command checks the status of an Indexer Cluster from the Cluster Master?

- A. splunk apply shcluster-bundle
- B. splunk list forward-server
- C. splunk show cluster-status
- D. splunk show shcluster-status

Q8: In which file do you configure an Indexer to join a cluster as a peer node?

- A. outputs.conf
- B. inputs.conf
- C. props.conf
- D. server.conf

Q9: What happens if the search factor (SF) is set to 2 in an Indexer Cluster?

- A. Each Search Head connects to two Indexers
- B. Two users can query the same data simultaneously
- C. Searches are throttled to two per user
- D. Each piece of data has two searchable copies

Q10: Which of the following is a reason to implement a distributed search architecture in Splunk?

- A. To reduce CPU usage on forwarders
- B. To replace the need for dashboards
- C. To improve performance and scalability across large environments
- D. To monitor a single local file

## SPLK-1003 Splunk Indexes

The index is the primary logical and physical unit of data organization, determining storage on disk and access permissions. Proper index management balances search performance with storage costs and compliance. Administrators must plan for a 20-25% disk buffer beyond estimated index sizes to accommodate spikes and maintenance tasks. If an index is not explicitly defined for a data source, Splunk defaults to the **main** index, which can lead to disorganized storage.

### 1. Index Types

Splunk supports Event Indexes for unstructured logs and Metrics Indexes for time-series numerical data. Metrics indexes are highly optimized for performance and resource efficiency. Additionally, Summary Indexes are used to store precomputed, aggregated results from long-term reports using the **collect** command, allowing for high-speed dashboarding without re-scanning raw data.

### 2. Bucket Lifecycle

As data ages, it moves through a bucket lifecycle: Hot, Warm, Cold, and Frozen. Hot buckets are actively writable and stored in the **homePath** (shared with warm). Once closed, they become Warm buckets, which are searchable but not writable. Cold buckets are moved to the **coldPath**, often on slower storage, and eventually, data moves to the Frozen state where it is deleted or archived. Manually restored frozen data is placed in the **thawedPath** for searchability.

### 3. Managing Indexes

Index configurations in **indexes.conf** define storage paths and retention limits based on time (**frozenTimePeriodInSecs**) or total size (**maxTotalDataSizeMB**). The **splunk fsck** tool is used to diagnose and repair index corruption or metadata issues, while **currentDBSizeMB** helps forecast disk requirements. Administrators must monitor these metrics to ensure the system does not hit disk-full scenarios that would halt ingestion.

Efficient indexing is contingent on complying with organizational license constraints.

### 4. Splunk Indexes Practice Question

Q1: Which of the following is optimized for storing time-series numerical data in Splunk?

- A. Event Index
- B. Summary Index
- C. Frozen Index
- D. Metrics Index

Q2: What is the default index used by Splunk if no index is explicitly specified?

- A. `_internal`
- B. `default`
- C. `main`
- D. `metrics`

Q3: Which configuration parameter in `indexes.conf` sets the maximum disk usage allowed for an index?

- A. `frozenTimePeriodInSecs`
- B. `thawedPath`
- C. `maxTotalDataSizeMB`
- D. `homePath`

Q4: What is the correct lifecycle transition for a bucket in Splunk?

- A. Cold → Warm → Hot → Thawed
- B. Hot → Warm → Cold → Frozen
- C. Frozen → Warm → Hot → Cold
- D. Hot → Cold → Warm → Frozen

Q5: Which index type is best for storing raw log events such as syslog or Apache logs?

- A. Metrics Index
- B. Lookup Index
- C. Summary Index
- D. Event Index

Q6: What happens to a frozen bucket by default in Splunk if no archiving is configured?

- A. It remains in coldPath indefinitely
- B. It is backed up to Splunk Cloud
- C. It is permanently deleted
- D. It is compressed and moved to thawedPath

Q7: What is the function of the `thawedPath` in `indexes.conf`?

- A. Controls license usage
- B. Stores buckets that are too large to index
- C. Stores event summaries
- D. Holds restored frozen data for search

Q8: Which SPL query retrieves index size, total event count, and retention settings?

- A. `| dbinspect index=*`
- B. `| rest /services/data/indexes`
- C. `| metadata type=indexes`
- D. `| tstats count from index=*`

Q9: What is the correct way to discard data after its retention period in Splunk?

- A. Configure `frozenTimePeriodInSecs`
- B. Use `btool` to delete buckets
- C. Move it to `_nullIndex`
- D. Disable all indexers

Q10: Which CLI command is used to create a new index in Splunk?

- A. `splunk add index`
- B. `splunk create index`
- C. `splunk define index`
- D. `splunk enable index`

## SPLK-1003 License Management

Splunk operates on a license model based on the raw volume of data ingested daily. Effective management ensures compliance with quotas while providing visibility into which sources consume the most capacity. In distributed environments, a License Master coordinates with license slaves, such as indexers, to monitor volume. Licensing is tracked at the indexing tier, and while forwarders are free, the data they transmit counts against the daily limit.

### 1. License Types and Roles

Several license types exist, including Enterprise, Free, Developer, and Trial licenses. The Trial license provides full Enterprise features for a limited eval period before reverting to the restricted Free version. The License

Master hosts and distributes **.lic** files to slaves and enforces daily indexing limits. Administrators can manage these roles via the License Manager in Splunk Web and back up the **\$SPLUNK\_HOME/etc/licenses/** directory for disaster recovery.

## 2. License Pools and Stacks

To manage consumption across departments, an organization's total license capacity is organized into a license stack, which can be subdivided into license pools. This logical division prevents a single high-volume source from consuming the entire capacity. The **splunk show license-status** command and the Monitoring Console provide visibility into pool-specific usage and historical trends.

## 3. Violations and Enforcement

Splunk tracks license violations in a rolling 30-day window. A violation occurs when daily ingestion exceeds the license limit. An organization is allowed a maximum of 5 violations in any 30-day period. On the 6th violation, search functionality is restricted for all non-admin users, although data indexing continues. The violation count cannot be cleared manually and only resets as the older violations roll off the 30-day window.

Before data is indexed and counted against a license, it must be properly parsed.

## 4. License Management Practice Question

Q1: What happens if the combined daily indexing volume across all license pools exceeds the license stack capacity?

- A. Only the largest pool is disabled
- B. Search functionality for admins is disabled
- C. Splunk deletes excess data
- D. A license violation occurs

Q2: Which Splunk license type provides full feature access and is intended for production use in large, distributed environments?

- A. Enterprise License
- B. Developer License
- C. Free License
- D. Trial License

Q3: Which license type is best suited for small-scale or personal Splunk deployments with limited daily indexing needs?

- A. Free License
- B. Developer License
- C. Trial License
- D. Enterprise License

Q4: In a distributed environment, which component acts as the centralized manager of license pools and indexing limits?

- A. Search Head
- B. License Master

- C. Deployment Server
- D. Monitoring Console

Q5: Which of the following best describes a license pool in Splunk?

- A. A backup location for license files
- B. A list of authorized users
- C. A logical division of license capacity with assigned ingestion limits
- D. A search result cache

Q6: What is the unit of measurement used by Splunk to enforce license usage?

- A. Compressed storage size
- B. Event count
- C. Raw data volume ingested before indexing
- D. Number of forwarders

Q7: Which of the following is a valid reason for setting up multiple license pools within an organization?

- A. Reducing data ingestion rate
- B. Enabling data forwarding
- C. Allocating daily indexing limits to specific departments
- D. Monitoring search performance

Q8: If your Splunk instance has exceeded the daily license volume limit for six days in a 30-day period, what happens?

- A. Data ingestion stops entirely
- B. All alerts are disabled
- C. Admin search capabilities are restricted
- D. Non-admin search functionality is disabled

Q9: Which command can be used from the CLI to check license status?

- A. splunk show stats
- B. splunk show license-status
- C. splunk view license-log
- D. splunk list license-indexers

Q10: What is the best way to proactively prevent license violations in a high-ingestion Splunk deployment?

- A. Reboot the Splunk instance daily
- B. Restrict access to the Search Head
- C. Set up alerts when usage nears the daily limit
- D. Disable indexing across all forwarders

The parsing phase transforms raw data into searchable events by tokenizing the stream, identifying timestamps, and assigning metadata. This phase is the "brain" of the ingestion pipeline and is where event boundaries are defined using **LINE\_BREAKER** rules. Decisions made here regarding index-time versus search-time extraction have long-term impacts on system flexibility and storage efficiency.

## 1. Parsing Overview

During parsing, Splunk performs tokenization, metadata assignment, and timestamp recognition. Tokenization breaks the raw stream into discrete events and tokens. Metadata assignment attaches Host, Source, and Sourcetype fields. Correct timestamp recognition is essential to place events accurately on the timeline, and administrators often use **props.conf** settings like **TIME\_FORMAT** to resolve non-standard logs.

## 2. Index-Time vs. Search-Time Extraction

Index-time parsing is rigid and increases storage volume but is necessary for event breaking, routing, and masking. The **INDEXED\_EXTRACTIONS** setting is used for structured data like CSV or JSON to improve search speed, though these fields become immutable once indexed. Most field extractions are deferred to search-time for flexibility, allowing administrators to modify definitions without re-indexing the raw data.

Parsing provides the opportunity to manipulate raw data for security and compliance.

## 3. Parsing Phase and Data Practice Question

Q1: What is the primary purpose of the parsing phase in Splunk?

- A. To store raw data in indexes
- B. To create dashboards for data visualization
- C. To forward data from one Splunk instance to another
- D. To transform raw data into structured, searchable events

Q2: Which configuration file defines timestamp recognition, line breaking, and field extraction behavior for a sourcetype?

- A. inputs.conf
- B. limits.conf
- C. indexes.conf
- D. props.conf

Q3: Which of the following props.conf settings helps Splunk identify the beginning of a timestamp in an event?

- A. TIME\_PREFIX
- B. KV\_MODE
- C. LINE\_BREAKER
- D. SHOULD\_LINEMERGE

Q4: What is the effect of setting 'SHOULD\_LINEMERGE = false' in props.conf?

- A. It instructs Splunk to break events based strictly on LINE\_BREAKER
- B. It enables default field extractions
- C. It disables parsing of timestamp fields
- D. It allows Splunk to automatically merge multiline events

Q5: Which transforms.conf setting is used to mask sensitive data such as credit card numbers?

- A. FORMAT = MetaData:Index
- B. FORMAT = nullQueue
- C. DEST\_KEY = \_raw
- D. REGEX = .\*

Q6: What is the purpose of the Field Extractor tool in Splunk?

- A. To define indexer cluster behavior
- B. To create and test field extraction rules graphically
- C. To adjust time zone recognition in props.conf
- D. To configure data forwarding

Q7: Which props.conf setting disables all field extractions during indexing and search?

- A. KV\_MODE = none
- B. FIELD\_EXTRACTION = disabled
- C. INDEXED\_EXTRACTIONS = none
- D. LINE\_BREAKER = none

Q8: How can you monitor queue usage to troubleshoot slow ingestion?

- A. Use netstat to check open connections
- B. Enable debug mode in inputs.conf
- C. Search \_internal for queue metrics
- D. Search \_audit index for latency

Q9: What is the role of DEST\_KEY = MetaData:Host in transforms.conf?

- A. Redirects data to a quarantine index
- B. Prevents the matched events from being indexed
- C. Drops field extractions during indexing
- D. Overrides the default host value of an event

Q10: Which internal index should be searched to debug parsing issues?

- A. \_metrics
- B. \_audit
- C. \_internal
- D. \_introspection

## SPLK-1003 Manipulating Raw Data

Data manipulation during ingestion is a strategic necessity for organizations complying with privacy regulations or optimizing storage. By transforming raw data before it is written to disk, administrators ensure that sensitive information is masked and only high-value events are retained. These manipulations are configured through the interplay of **props.conf** and **transforms.conf**.

## 1. Configuration Basics

Manipulating raw data involves using **props.conf** to identify the data source and **transforms.conf** to define the specific transformation rules using **REGEX** and **FORMAT**. This allows for the redaction of sensitive details like credit card numbers or the standardization of field names. Administrators should minimize regex complexity to avoid performance bottlenecks and document all rules for future audits.

## 2. Advanced Techniques

Advanced manipulation utilizes the **DEST\_KEY** parameter in **transforms.conf**. Setting **DEST\_KEY = \_raw** allows for masking information directly within the raw event data, while **\_MetaData:Index** is used for dynamic routing, such as sending "ERROR" logs to a specific index. These index-time transformations are irreversible and must be tested thoroughly in a staging environment.

## 3. TRANSFORMS vs. REPORT

It is critical to distinguish between the **TRANSFORMS** and **REPORT** stanzas in **props.conf**. **TRANSFORMS** rules are applied during the parsing phase at index-time and can modify **\_raw** data or metadata. **REPORT** rules define search-time extractions that only affect how data is displayed in search results. If a task requires modifying how data is indexed or stored, the **TRANSFORMS** attribute must be used.

Fine-tuning these manipulations ensures optimal performance and reduced noise in the system.

## 4. Manipulating Raw Data Practice Question

Q1: Which setting in transforms.conf is responsible for masking sensitive data in the raw event?

- A. FORMAT = nullQueue
- B. SOURCE\_KEY = \_meta
- C. DEST\_KEY = \_raw
- D. REGEX = .P//.

Q2: In which configuration file do you associate a sourcetype with a transformation rule?

- A. props.conf
- B. inputs.conf
- C. limits.conf
- D. server.conf

Q3: What is the purpose of setting DEST\_KEY = \_MetaData:Index in a transform?

- A. To enrich data with external lookups
- B. To discard unwanted events
- C. To mask sensitive fields
- D. To reroute events to a specific index

Q4: How can you route logs containing "ERROR" to a separate index?

- A. Set KV\_MODE = none in props.conf
- B. Create a new index in indexes.conf

- C. Use LOOKUP in transforms.conf
- D. Define a regex-based routing rule using `DEST_KEY = _MetaData:Index`

Q5: If you want to rename the field `http_user_agent` to `user_agent`, which setting should you include in transforms.conf?

- A. `SOURCE_KEY = http_user_agent`
- B. `DEST_KEY = nullQueue`
- C. `REGEX = .*`
- D. `FORMAT = sourcetype::user_agent`

Q6: What is the role of `FORMAT = nullQueue` in a transformation rule?

- A. Masks sensitive information in the event
- B. Drops matched events before indexing
- C. Routes events to a quarantine index
- D. Changes the sourcetype of the event

Q7: What is the benefit of using modular configurations for transformations?

- A. Reduces event breaking
- B. Ensures all fields are automatically extracted
- C. Simplifies management and scalability
- D. Forces routing to a default index

Q8: To add a static field `environment=production` to logs from prod servers, which key should be used in transforms.conf?

- A. `DEST_KEY = _meta`
- B. `DEST_KEY = _raw`
- C. `DEST_KEY = queue`
- D. `DEST_KEY = _MetaData:Source`

Q9: What does the following regex in transforms.conf extract: `REGEX =`

`"product" : "(?P<product>[^\"]+)"?`

- A. It drops all events except those with product info
- B. It extracts the product field from a nested JSON
- C. It masks the product value from the event
- D. It routes the product logs to a separate index

Q10: Which method can help reduce CPU usage during transformation?

- A. Enabling event merging in props.conf
- B. Using efficient regular expressions with non-greedy quantifiers
- C. Increasing sourcetype complexity
- D. Applying transformations globally to all inputs

# SPLK-1003 Fine Tuning Inputs

Input optimization ensures system stability by reducing noise and improving data reliability. By fine-tuning configurations, administrators prevent indexing delays and duplicated events. This involves a continuous process of auditing inputs and applying resource-efficient settings to balance the needs for performance and data durability.

## 1. Optimization Techniques

Key techniques include throttling data ingestion rates via polling intervals and applying filters to exclude low-value "noise." Regex filters in **props.conf** and **transforms.conf** can drop repetitive debug logs before they reach the indexer, saving license capacity and storage. Additionally, disabling unnecessary automatic field extractions in **props.conf** can reduce the processing load during ingestion.

## 2. Reliability vs. Performance

Administrators must balance data durability with system throughput. Enabling Indexer Acknowledgment (**useACK = true**) and data compression (**compressed = true**) ensures reliable delivery across the network but adds latency and CPU overhead. Choosing the right protocol, such as TCP for mission-critical logs versus UDP for non-critical telemetry, is a fundamental part of this tuning process.

All configuration changes must be validated in a controlled environment before production.

## 3. Fine Tuning Inputs Practice Question

Q1: What does setting 'interval = 300' in inputs.conf achieve?

- A. It disables the input after 300 seconds
- B. It sets the ingestion rate to 300 events per second
- C. It changes the default index to 300
- D. It instructs Splunk to poll the file every 5 minutes

Q2: Which stanza in transforms.conf is used to exclude events based on pattern matching?

- A. exclude\_debug\_events
- B. queueThrottle
- C. lineBreaker
- D. eventIndexer

Q3: What does setting 'KV\_MODE = none' in props.conf do?

- A. Sets the key-value separator to none
- B. Enables automatic field extraction
- C. Combines multiple events into one
- D. Disables automatic key-value field extraction

Q4: How can you ensure that Splunk does not ingest log lines containing 'DEBUG'?

- A. Set interval = 0
- B. Define a regex in transforms.conf with FORMAT = nullQueue

- C. Set sourcetype = DEBUG
- D. Disable the monitor stanza

Q5: Which configuration file is used to throttle data input frequency?

- A. indexes.conf
- B. transforms.conf
- C. inputs.conf
- D. props.conf

Q6: What is the effect of assigning static metadata in inputs.conf?

- A. Increases indexing throughput
- B. Reduces processing overhead at ingest time
- C. Disables indexing
- D. Prevents field extraction

Q7: Which props.conf setting disables all field extractions during indexing and search?

- A. KV\_MODE = none
- B. FIELD\_EXTRACTION = disabled
- C. LINE\_BREAKER = none
- D. INDEXED\_EXTRACTIONS = none

Q8: How can you monitor queue usage to troubleshoot slow ingestion?

- A. Search `_audit` index for latency
- B. Use netstat to check open connections
- C. Enable debug mode in inputs.conf
- D. Search `_internal` for queue metrics

Q9: What is the result of configuring `FORMAT = nullQueue` in transforms.conf?

- A. Drops field extractions during indexing
- B. Prevents the matched events from being indexed
- C. Redirects data to a quarantine index
- D. Applies compression to forwarded data

Q10: Which of the following would help in fine-tuning Splunk performance at the input stage?

- A. Indexing all debug and trace logs
- B. Increasing line merging
- C. Reducing field extraction and enabling throttling
- D. Adding field extractions at search time

## SPLK-1003 Getting Data In – Staging

The staging environment is a critical risk mitigation tool that allows administrators to validate ingestion configurations without impacting production data. It provides a safe space to test complex parsing rules, event

breaking for multi-line logs, and metadata assignments. Configurations validated in staging are promoted to production with high confidence in their accuracy.

## 1. Staging Environment Overview

A staging environment consists of a standalone Splunk instance mirroring production settings. Administrators ingest representative data samples to verify that timestamps are correctly parsed and that **LINE\_BREAKER** settings properly handle multi-line events. Using version control for **.conf** files during this phase ensures that all changes are documented and can be audited before promotion.

## 2. Validation Checks

Administrators perform validation checks using the **splunk add oneshot** command to ingest a file once for format validation, or the **monitor** command to simulate continuous ingestion. The **btprobe** utility is also used to debug bucket and timestamp parsing issues. Once data appears correctly in the staging search interface, the validated configurations are promoted to the production deployment server.

Once data is correctly ingested, managing the users who access that data becomes the priority.

## 3. Getting Data In – Staging Practice Question

Q1: What is the primary purpose of a staging environment in Splunk data onboarding?

- A. To archive old logs
- B. To monitor user search activity
- C. To test and validate inputs before deploying to production
- D. To store high-volume real-time data

Q2: Which command is used to ingest a file into Splunk one time for validation?

- A. splunk import file
- B. splunk add oneshot
- C. splunk start input
- D. splunk add monitor

Q3: What file should you edit to validate timestamp extraction rules in staging?

- A. limits.conf
- B. server.conf
- C. props.conf
- D. transforms.conf

Q4: How can you simulate production data volume in a staging environment?

- A. Restart Splunk repeatedly
- B. Generate test logs using a script and ingest them
- C. Disable indexing to speed up ingestion
- D. Use dashboards to visualize traffic

Q5: Which SPL command helps validate if the timestamp parsing is accurate?

- A. `index=staging_index | stats count by _time`

AAAdemy | <https://www.aaademy.com>

- B. index=\_internal | top sourcetype
- C. index=staging\_index | table host
- D. index=\_audit | stats count

Q6: Where is the default internal log source used for monitoring indexing throughput?

- A. \_metrics
- B. \_audit
- C. \_internal
- D. \_introspection

Q7: Which file allows you to assign metadata like **host** and **sourcetype** for file inputs?

- A. limits.conf
- B. web.conf
- C. fields.conf
- D. inputs.conf

Q8: What should be used to test if field extractions are working as expected in staging?

- A. a dashboard
- B. server.conf
- C. index=\_audit
- D. a search using extracted fields

Q9: If metadata appears incorrect in events, which action should be taken first?

- A. Delete the data and re-ingest it
- B. Restart the indexer
- C. Review inputs.conf and props.conf
- D. Modify server.conf

Q10: Which command can be used to debug active input settings in Splunk?

- A. splunk list metadata
- B. splunk btool inputs list --debug
- C. splunk check config
- D. splunk list inputs

## SPLK-1003 Splunk User Management

Effective user administration is grounded in the principle of least privilege, ensuring users have access only to necessary features. Splunk provides predefined roles—Admin, Power, and User—as well as specialized roles like **can\_delete**, which grants the ability to permanently remove data using the **delete** command. The **splunk\_system\_user** is a reserved internal role that should never be assigned to human users.

### 1. User Roles and Capabilities

Capabilities define specific actions a role can perform, such as scheduling searches or managing other users. The Admin role has full system control, while the User role is typically restricted to basic searching. By carefully assigning these capabilities, administrators control the functional scope of every user group. The **can\_delete** role is particularly sensitive and is not enabled by default.

## 2. Custom Roles and Access Control

Custom roles are created in **authorize.conf** to meet specific organizational needs, often inheriting capabilities from predefined roles. The most common use for custom roles is index-level access control, restricting a role to searching only specific indexes. This ensures data sovereignty and prevents unauthorized access to sensitive information while maintaining a centralized search capability.

Authentication mechanisms provide the secure gateway for these managed users.

## 3. Splunk User Management Practice Question

Q1: Which predefined role in Splunk grants full administrative access to the system?

- A. Can\_delete
- B. Power
- C. User
- D. Admin

Q2: What capability allows a user role to create and edit dashboards in Splunk?

- A. edit\_user
- B. run\_search
- C. edit\_dashboard
- D. schedule\_search

Q3: Which of the following is NOT a predefined role in Splunk?

- A. Power
- B. Admin
- C. User
- D. Viewer

Q4: Which file can be used to manage users manually in Splunk?

- A. authentication.conf
- B. inputs.conf
- C. authorize.conf
- D. web.conf

Q5: Where in Splunk Web do you go to assign a user to a role?

- A. Settings > Roles > Permissions
- B. Settings > Users and Authentication > Users
- C. Monitoring Console > Users
- D. Search & Reporting > Permissions

Q6: What is the purpose of a search filter (`srchFilter`) in a role configuration?

- A. To hide fields in dashboards
- B. To restrict data users can search
- C. To schedule user activity
- D. To filter incoming data

Q7: If a user cannot create a dashboard, what is the most likely missing capability?

- A. `edit_user`
- B. `edit_alerts`
- C. `edit_dashboard`
- D. `schedule_search`

Q8: Which of the following describes the 'Power' role?

- A. Can create reports and dashboards
- B. Has access to internal logs only
- C. Can delete indexed data
- D. Can manage users and roles

Q9: What is the best method to assign access to only the `finance_logs` index?

- A. Set the index in a dashboard token
- B. Assign index access in a custom role
- C. Use a field alias
- D. Create a search macro

Q10: Which index contains user activity logs such as login events and role changes?

- A. `_metrics`
- B. `_internal`
- C. `_audit`
- D. `_introspection`

## SPLK-1003 Splunk Authentication Management

Securing access to Splunk involves diverse authentication backends, ranging from built-in accounts to enterprise identity providers. By integrating Splunk with existing systems, administrators streamline access while maintaining security standards. It is critical to distinguish between authentication (verifying identity) and authorization (determining permissions via roles).

### 1. Authentication Methods

Splunk supports built-in authentication, LDAP integration, and Single Sign-On (SSO) via SAML. Built-in authentication is suitable for small setups, while LDAP and SAML are the standards for enterprises. Multi-Factor Authentication (MFA) is strongly recommended for admin roles and is typically implemented through the SSO

provider or a reverse proxy. Administrators can use **splunk list auth-tokens** to view active session tokens and troubleshoot SSO handshakes.

## 2. Configuration and Troubleshooting

Authentication settings are managed in **authentication.conf**. Common troubleshooting involve checking for clock skew between the Identity Provider and Splunk, or verifying the **userNameAttribute** in LDAP configurations to ensure it matches the attribute used by the directory service. The **\_audit** index is a vital resource for monitoring login activity and identifying failed authentication attempts.

The entire administration framework is governed by the underlying configuration file architecture.

## 3. Splunk Authentication Management Practice Question

Q1: Which of the following authentication methods allows Splunk to integrate with Active Directory for centralized user management?

- A. Built-in Authentication
- B. SAML
- C. LDAP
- D. OAuth

Q2: Where are user credentials stored when using Splunk's built-in authentication method?

- A. authorize.conf
- B. authentication.conf
- C. users.conf
- D. web.conf

Q3: In a SAML-based SSO configuration for Splunk, which of the following is required?

- A. An LDAP server IP
- B. A field extraction in props.conf
- C. A role defined in transforms.conf
- D. An Entity ID and a certificate

Q4: What is the primary benefit of enabling Single Sign-On (SSO) in Splunk environments?

- A. Seamless access using centralized credentials
- B. Automatic field extractions for dashboards
- C. Enhanced password recovery tools
- D. Better data parsing performance

Q5: Which CLI command can be used to create a new user with a specified role in Splunk?

- A. splunk add role --user
- B. splunk create-user
- C. splunk add user
- D. splunk create user

Q6: In the LDAP configuration for Splunk, which stanza maps external groups to internal roles?

- A. [authentication\_roles]

- B. [roleMap\_LDAP]
- C. [ldap\_roles]
- D. [ldap\_group\_mapping]

Q7: Which Splunk index stores login attempts and authentication-related audit logs?

- A. \_audit
- B. \_introspection
- C. \_internal
- D. \_metrics

Q8: Which of the following statements is true about Multi-Factor Authentication (MFA) in Splunk?

- A. MFA disables role-based access
- B. MFA is configured directly in Splunk's built-in UI
- C. MFA can only be implemented through an SSO provider
- D. MFA is applied by editing transforms.conf

Q9: What does the `SSLEnabled = 1` setting do in an LDAP authentication stanza?

- A. It disables password caching
- B. It encrypts the LDAP traffic between Splunk and the LDAP server
- C. It forces users to change passwords
- D. It enables two-factor login

Q10: What is the recommended way to verify that an LDAP user was correctly mapped to a Splunk role?

- A. Restart the Splunk service
- B. Search the \_internal index for search logs
- C. Check the LDAP server directly
- D. Attempt login and review role assignment in Splunk Web

## SPLK-1003 Splunk Configuration Files

Splunk is a configuration-driven platform where system behavior is defined by **.conf** files. Mastering the purpose and management of these files is the most critical skill for an administrator. Understanding how settings are applied through layers and precedence allows for precise control over the environment.

### 1. Critical Configuration Files

The primary files governing the system include **inputs.conf** for collection, **outputs.conf** for forwarding, and **props.conf** and **transforms.conf** for parsing. Additionally, **fields.conf** defines whether a field is indexed or searchable and its multi-value behavior, while **limits.conf** defines system performance thresholds like **maxKBps**. These files form the foundation of every administrative task in Splunk.

### 2. Maintenance and Debugging

Effective maintenance requires an understanding of configuration reload behaviors. While many changes to **props.conf** can be hot-reloaded using the **splunk \_internal call /services/admin/config-roll** command, changes to **inputs.conf** or **indexes.conf** typically require a service restart. Throughout all tasks, the **splunk btool** utility remains the essential tool for debugging, providing a final view of the effective configuration to ensure system integrity.

Maintaining a production-ready Splunk environment requires a commitment to architectural best practices, proactive monitoring through the Monitoring Console, and the rigorous validation of configurations in staging. By adhering to the principles of least privilege and optimizing ingestion through filtering and throttling, administrators ensure that Splunk remains a reliable tool for operational intelligence. Regular audits of license usage, forwarder health, and user access are necessary to sustain the long-term health and security of the platform.

### 3. Splunk Configuration Files Practice Question

Q1: Which configuration file is used to define how data should be parsed, including timestamp formats and field extractions?

- A. inputs.conf
- B. outputs.conf
- C. props.conf
- D. fields.conf

Q2: In transforms.conf, what does setting **DEST\_KEY = queue** and **FORMAT = nullQueue** accomplish?

- A. It sends data to a specific index
- B. It encrypts the raw data
- C. It masks field names
- D. It discards the event entirely

Q3: Which configuration file is responsible for field extraction rules referenced by the **REPORT** stanza in props.conf?

- A. inputs.conf
- B. outputs.conf
- C. transforms.conf
- D. limits.conf

Q4: Which configuration file defines the data sources Splunk will ingest from files, directories, or network ports?

- A. inputs.conf
- B. outputs.conf
- C. transforms.conf
- D. limits.conf

Q5: What does **autoLB = true** enable in an outputs.conf configuration?

- A. Authentication for data forwarding
- B. Load balancing across multiple Indexers
- C. Automatic port scanning
- D. Data deduplication at index time

Q6: What is the effect of using `BREAK_ONLY_BEFORE` in `props.conf`?

- A. Replaces matching field values
- B. Filters out malformed events
- C. Splits events based on a regex pattern
- D. Disables multi-line merging

Q7: Which of the following causes events matching a certain REGEX to be excluded from indexing in `transforms.conf`?

- A. `FORMAT = masked_field`
- B. `FORMAT = metadataOverride`
- C. `FORMAT = routeToIndex`
- D. `FORMAT = nullQueue`

Q8: What is the effect of placing configuration changes in the `local/` directory of a Splunk app?

- A. They override default configurations
- B. They are ignored unless signed
- C. They are less secure than system-wide configs
- D. They disable field extraction

Q9: What is the purpose of the `btool` command in Splunk?

- A. To query indexed data
- B. To troubleshoot and validate configuration file values
- C. To search license usage
- D. To forward logs to Indexers

Q10: Which of the following directives in `props.conf` defines how timestamp formats are handled during data parsing?

- A. `TRANSFORMS-routing`
- B. `BREAK_ONLY_BEFORE`
- C. `TIME_FORMAT`
- D. `MAX_DAYS_AGO`

## Learning Path & Study Advice

A strong preparation path should begin with the administrative fundamentals of Splunk Enterprise, especially platform structure, configuration behavior, and the role of indexes, users, and licensing in a working deployment. From there, study should move into data ingestion topics in a logical progression: general data onboarding, staging practices, forwarder configuration, input methods, and the parsing pipeline. Distributed search should be approached after the learner is comfortable with core standalone administration, since it builds on foundational understanding and introduces architectural coordination across components.

Candidates benefit most when they study by connecting each knowledge area to the full data lifecycle. Rather than treating topics as isolated features, it is more effective to understand how configuration decisions influence ingestion, how ingestion affects parsing, how parsing affects indexing, and how indexing affects search and operational outcomes. Practical comprehension is especially important in areas such as forwarder behavior, authentication design, index planning, and input tuning. A good study approach emphasizes cause-and-effect relationships, configuration logic, and operational reasoning over memorization.

## Who This PDF Is For

This PDF is intended for system administrators, platform administrators, security operations personnel, and IT professionals who work with or support Splunk Enterprise environments. It is most suitable for learners who already have basic familiarity with Splunk and want to develop a stronger administrative understanding of how the platform is configured, secured, and maintained. It will be particularly useful for those preparing for an administration-focused certification path, as well as for professionals who need a structured overview of the key knowledge areas involved in managing Splunk Enterprise in practice.

## Call To Action

This document provides an overview of structured learning and certification preparation approaches. For learners seeking clear knowledge organization, guided study planning, and exam-focused practice resources, AAAdemy offers a comprehensive platform to support independent and effective learning.

Explore additional training materials, study guidance, and practice resources at:

[Splunk SPLK-1003 Enterprise Certified Admin Certification Training Course - AAAdemy](#)

Online Flashcards (Quizlet):

<https://quizlet.com/user/AAAdemy/folders/splk-1003-splunk-enterprise-certified-admin?i=6zfa5t&x=1xqt>

### Splunk Admin Basics Practice Question

A1: Answer: C

Explanation: The Search Head handles search management, dashboards, and alerts, but does not store indexed data.

A2: Answer: B

Explanation: The Deployment Server is used to distribute configuration apps to forwarders.

A3: Answer: D

Explanation: The Universal Forwarder is lightweight and designed for efficient data forwarding.

A4: Answer: B

Explanation: The Indexer is responsible for storing, indexing, and making data searchable.

A5: Answer: D

Explanation: Port 8000 is the default port for accessing Splunk Web.

A6: Answer: D

Explanation: The parsing stage handles timestamp assignment and event breaking.

A7: Answer: A

Explanation: The btool command is used to debug and validate configuration files.

A8: Answer: A

Explanation: The indexing stage writes processed data into index buckets.

A9: Answer: C

Explanation: splunkd.log contains general system and error logs.

A10: Answer: A

Explanation: Heavy Forwarders can parse and transform data before forwarding.

### License Management Practice Question

A1: Answer: D

Explanation: Exceeding license capacity triggers a license violation.

A2: Answer: A

Explanation: Enterprise License is designed for full-feature production environments.

A3: Answer: A

Explanation: Free License is suitable for small-scale usage.

A4: Answer: B

Explanation: The License Master manages license pools and usage.

A5: Answer: C

Explanation: License pools divide indexing capacity logically.

A6: Answer: C

Explanation: Licensing is based on raw data volume ingested.

A7: Answer: C

Explanation: Multiple pools allow allocation by department.

A8: Answer: D

Explanation: Non-admin searches are disabled after repeated violations.

A9: Answer: B

Explanation: splunk show license-status checks license usage.

A10: Answer: C

Explanation: Alerts help prevent exceeding license limits.

#### Splunk Configuration Files Practice Question

A1: Answer: C

Explanation: props.conf defines parsing rules and field extraction.

A2: Answer: D

Explanation: nullQueue discards matching events.

A3: Answer: C

Explanation: transforms.conf contains field extraction rules.

A4: Answer: A

Explanation: inputs.conf defines data inputs.

A5: Answer: B

Explanation: autoLB enables load balancing across indexers.

A6: Answer: C

Explanation: BREAK\_ONLY\_BEFORE splits events using regex.

A7: Answer: D

Explanation: nullQueue discards events from indexing.

A8: Answer: A

Explanation: local/ overrides default configurations.

A9: Answer: B

Explanation: btool is used for config troubleshooting.

A10: Answer: C

Explanation: TIME\_FORMAT defines timestamp parsing.

#### Splunk Indexes Practice Question

A1: Answer: D

Explanation: Metrics Index is optimized for time-series numeric data.

A2: Answer: C

Explanation: main is the default index.

A3: Answer: C

Explanation: maxTotalDataSizeMB controls index size.

A4: Answer: B

Explanation: Correct lifecycle: Hot → Warm → Cold → Frozen.

A5: Answer: D

Explanation: Event Index stores raw log data.

A6: Answer: C

Explanation: Frozen buckets are deleted if not archived.

A7: Answer: D

Explanation: thawedPath stores restored frozen data.

A8: Answer: B

Explanation: REST endpoint retrieves index details.

A9: Answer: A

Explanation: frozenTimePeriodInSecs controls retention.

A10: Answer: A

Explanation: splunk add index creates a new index.

### Splunk User Management Practice Question

A1: Answer: D

Explanation: Admin role has full system access.

A2: Answer: C

Explanation: edit\_dashboard allows dashboard creation/editing.

A3: Answer: D

Explanation: Viewer is not a predefined role.

A4: Answer: A

Explanation: authentication.conf manages authentication settings.

A5: Answer: B

Explanation: Users are managed under Users and Authentication.

A6: Answer: B

Explanation: srchFilter restricts searchable data.

A7: Answer: C

Explanation: edit\_dashboard is required to create dashboards.

A8: Answer: A

Explanation: Power role can create reports and dashboards.

A9: Answer: B

Explanation: Index access is controlled via roles.

A10: Answer: C

Explanation: `_audit` logs user activity.

Splunk Authentication Management Practice Question

A1: Answer: C

Explanation: LDAP integrates with Active Directory.

A2: Answer: C

Explanation: `users.conf` stores local credentials.

A3: Answer: D

Explanation: SAML requires entity ID and certificate.

A4: Answer: A

Explanation: SSO provides centralized authentication.

A5: Answer: C

Explanation: `splunk add user` creates a new user.

A6: Answer: B

Explanation: `roleMap_LDAP` maps groups to roles.

A7: Answer: A

Explanation: `_audit` stores authentication logs.

A8: Answer: C

Explanation: MFA is implemented via external providers.

A9: Answer: B

Explanation: `SSLEnabled` encrypts LDAP communication.

A10: Answer: D

Explanation: Attempt login and review role assignment in Splunk Web.

Getting Data In Practice Question

A1: Answer: A

Explanation: `inputs.conf` defines data inputs.

A2: Answer: C

Explanation: `sourcetype` controls parsing behavior.

A3: Answer: C

Explanation: `sourcetype` defines data format.

A4: Answer: D

Explanation: Network Input collects UDP data.

A5: Answer: B

Explanation: HEC is best for JSON ingestion.

A6: Answer: B

Explanation: recursive enables subdirectory monitoring.

A7: Answer: D

Explanation: tcp:// defines TCP input.

A8: Answer: A

Explanation: host identifies source system.

A9: Answer: C

Explanation: metrics.log tracks throughput.

A10: Answer: C

Explanation: HEC receives HTTP event data.

#### Distributed Search Practice Question

A1: Answer: C

Explanation: The Search Head manages and executes search queries in a distributed environment.

A2: Answer: C

Explanation: Replication Factor determines how many copies of data are stored.

A3: Answer: B

Explanation: distsearch.conf is used to configure distributed search connections.

A4: Answer: A

Explanation: The Deployer distributes configuration bundles to Search Head Cluster members.

A5: Answer: B

Explanation: max\_searches\_per\_cpu controls concurrent search execution.

A6: Answer: A

Explanation: Search Head clustering provides high availability and load balancing.

A7: Answer: C

Explanation: splunk show cluster-status checks indexer cluster health.

A8: Answer: D

Explanation: server.conf is used for indexer cluster configuration.

A9: Answer: D

Explanation: Search Factor defines searchable copies of data.

A10: Answer: C

Explanation: Distributed search improves scalability and performance.

#### Getting Data In – Staging Practice Question

A1: Answer: C

Explanation: Staging environments are used to test configurations before production deployment.

A2: Answer: B

Explanation: splunk add oneshot ingests a file once for testing.

A3: Answer: C

Explanation: props.conf controls timestamp parsing rules.

A4: Answer: B

Explanation: Generating test logs simulates production load.

A5: Answer: A

Explanation: stats by \_time validates timestamp parsing.

A6: Answer: C

Explanation: \_internal index stores ingestion metrics.

A7: Answer: D

Explanation: inputs.conf assigns metadata like host and sourcetype.

A8: Answer: D

Explanation: Searching extracted fields verifies correctness.

A9: Answer: C

Explanation: inputs.conf and props.conf define metadata behavior.

A10: Answer: B

Explanation: btool helps debug input configurations.

#### Configuring Forwarders Practice Question

A1: Answer: A

Explanation: Universal Forwarders collect and forward data only.

A2: Answer: C

Explanation: outputs.conf defines indexer destinations.

A3: Answer: D

Explanation: Heavy Forwarders can parse and transform data.

A4: Answer: B

Explanation: deploymentclient.conf connects forwarders to Deployment Server.

A5: Answer: C

Explanation: autoLB enables load balancing.

A6: Answer: B

Explanation: splunk list forward-server shows connection status.

A7: Answer: D

Explanation: inputs.conf defines monitored inputs.

A8: Answer: C

Explanation: outputs.conf configures SSL forwarding.

A9: Answer: A

Explanation: Deployment Server centrally manages forwarders.

A10: Answer: C

Explanation: btool verifies active input configuration.

#### Forwarder Management Practice Question

A1: Answer: C

Explanation: Deployment Server manages forwarder configurations.

A2: Answer: B

Explanation: deploymentclient.conf defines DS connection.

A3: Answer: B

Explanation: metrics.log tracks forwarder throughput.

A4: Answer: B

Explanation: Server classes group forwarders.

A5: Answer: D

Explanation: outputs.conf misconfiguration prevents forwarding.

A6: Answer: C

Explanation: splunk show deploy-clients lists connected clients.

A7: Answer: B

Explanation: autoLB enables load balancing.

A8: Answer: C

Explanation: Reducing monitored inputs lowers CPU usage.

A9: Answer: D

Explanation: Compression reduces bandwidth usage.

A10: Answer: D

Explanation: Deployment apps should be centralized and backed up.

#### Monitor Inputs Practice Question

A1: Answer: C

Explanation: monitor:// ingests files and directories.

A2: Answer: B

Explanation: whitelist includes matching files.

A3: Answer: D

Explanation: `_internal` index stores logs and metrics.

A4: Answer: D

Explanation: Missing `crcSalt` may cause file skipping.

A5: Answer: C

Explanation: `blacklist = debug.log$` excludes that file.

A6: Answer: A

Explanation: `followTail` ingests only new data.

A7: Answer: A

Explanation: `whitelist + blacklist` controls inclusion/exclusion.

A8: Answer: C

Explanation: `metrics.log` shows ingestion throughput.

A9: Answer: B

Explanation: `recursive` enables subdirectory monitoring.

A10: Answer: A

Explanation: `interval` controls polling frequency.

#### Network and Scripted Inputs Practice Question

A1: Answer: C

Explanation: `interval` controls how often scripts run.

A2: Answer: C

Explanation: Syslog servers improve reliability and centralization.

A3: Answer: A

Explanation: HEC logs appear in `_internal` index.

A4: Answer: C

Explanation: SSL/TLS secures HEC endpoints.

A5: Answer: C

Explanation: Default HEC port is 8088.

A6: Answer: D

Explanation: `_internal` index contains HEC logs.

A7: Answer: B

Explanation: Scripts must have execution permissions.

A8: Answer: A

Explanation: Scripts go in \$SPLUNK\_HOME/bin/scripts/.

A9: Answer: B

Explanation: TCP provides reliable delivery.

A10: Answer: D

Explanation: script:// defines scripted input.

#### Agentless Inputs Practice Question

A1: Answer: D

Explanation: WMI allows agentless collection of Windows logs.

A2: Answer: A

Explanation: HTTPS ensures secure transmission of API data.

A3: Answer: B

Explanation: WMI can consume significant system resources.

A4: Answer: C

Explanation: Port 135 is used for RPC endpoint mapper.

A5: Answer: B

Explanation: Scripted inputs allow periodic API data collection.

A6: Answer: C

Explanation: inputs.conf defines scripted input settings.

A7: Answer: C

Explanation: Batching and retry logic mitigate API rate limits.

A8: Answer: B

Explanation: Credentials should be stored securely (env vars or encrypted).

A9: Answer: B

Explanation: Running the script manually helps debug issues.

A10: Answer: A

Explanation: WMI provides lightweight agentless data collection.

#### Fine Tuning Inputs Practice Question

A1: Answer: D

Explanation: interval = 300 means polling every 5 minutes.

A2: Answer: A

Explanation: Custom transform stanzas can exclude events.

A3: Answer: D

Explanation: KV\_MODE = none disables automatic KV extraction.

A4: Answer: B

Explanation: nullQueue drops matching events.

A5: Answer: C

Explanation: inputs.conf controls ingestion interval.

A6: Answer: B

Explanation: Static metadata reduces processing overhead.

A7: Answer: A

Explanation: KV\_MODE = none disables field extraction.

A8: Answer: D

Explanation: \_internal index contains queue metrics.

A9: Answer: B

Explanation: nullQueue prevents indexing of events.

A10: Answer: C

Explanation: Reducing extraction and throttling improves performance.

#### Parsing Phase and Data Practice Question

A1: Answer: D

Explanation: Parsing converts raw data into structured events.

A2: Answer: D

Explanation: props.conf defines parsing behavior.

A3: Answer: A

Explanation: TIME\_PREFIX identifies timestamp start.

A4: Answer: A

Explanation: SHOULD\_LINEMERGE = false enforces LINE\_BREAKER.

A5: Answer: C

Explanation: DEST\_KEY = \_raw enables masking.

A6: Answer: B

Explanation: Field Extractor provides GUI-based extraction.

A7: Answer: A

Explanation: KV\_MODE = none disables extraction.

A8: Answer: C

Explanation: \_internal index shows queue metrics.

A9: Answer: D

Explanation: MetaData:Host overrides host value.

A10: Answer: C

Explanation: `_internal` index is used for parsing/debug logs.

#### Manipulating Raw Data Practice Question

A1: Answer: C

Explanation: `DEST_KEY = _raw` modifies raw event content.

A2: Answer: A

Explanation: `props.conf` links sourcetypes to transforms.

A3: Answer: D

Explanation: `MetaData:Index` reroutes events to a different index.

A4: Answer: D

Explanation: `Regex + MetaData:Index` routes ERROR logs.

A5: Answer: A

Explanation: `SOURCE_KEY` identifies the field to modify.

A6: Answer: B

Explanation: `nullQueue` drops events before indexing.

A7: Answer: C

Explanation: Modular configs improve scalability and manageability.

A8: Answer: A

Explanation: `_meta` is used to add static metadata fields.

A9: Answer: B

Explanation: `Regex` extracts JSON field values.

A10: Answer: B

Explanation: Efficient `regex` reduces CPU usage.