SPLK-3001 Training Course

Splunk Enterprise Security Certified Admin

Structured Learning & Certification Preparation

Table of Contents

Introduction

The SPLK-3001 certification is designed to validate administrative knowledge and operational capability within Splunk Enterprise Security. It represents the ability to work with a security analytics platform that supports threat detection, investigation, contextual enrichment, and security operations workflows. In current enterprise environments, where organizations rely on centralized visibility and timely response to security events, this certification is relevant for professionals responsible for maintaining and improving a SIEM-driven security program.

About This Training / Certification

This certification is best understood as an intermediate-level credential focused on the administration and practical use of Splunk Enterprise Security. It assesses whether a candidate can understand the platform's core security workflows, deploy and configure essential components, validate the quality of ingested data, and support the content and context needed for effective detection and investigation. Within a broader learning journey, it usually builds on prior familiarity with Splunk fundamentals and supports progression into security operations, content engineering, detection management, and platform administration responsibilities.

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

Domain Area: ES Introduction

Candidates are expected to understand the purpose and structure of Splunk Enterprise Security as a specialized security solution built on top of Splunk. This includes understanding its role in supporting security monitoring, investigation, and contextual analysis across large volumes of machine and event data. A strong conceptual foundation in how Enterprise Security differs from general log search and reporting is important for understanding the rest of the platform.

Domain Area: Monitoring and Investigation

This area focuses on how analysts and administrators interact with security events inside Enterprise Security. Candidates should understand the workflows used to monitor notable events, review alert activity, triage suspicious behavior, and investigate incidents using available context and linked data. The emphasis is on understanding how the platform supports operational visibility and structured investigation rather than simply generating alerts.

Domain Area: Security Intelligence

This area covers the broader intelligence layer within Enterprise Security, including how security-relevant information is organized, interpreted, and used to improve visibility. Candidates should understand how the platform helps transform raw event data into meaningful security insights through correlation, context, prioritization, and operational interpretation.

Domain Area: Forensics, Glass Tables, and Navigation Control

Candidates should understand specialized interfaces and investigative views used within Enterprise Security. This includes the purpose of forensic-style analysis features, visual security dashboards such as glass tables, and the navigation structures that help users move between security content and workflows. The conceptual goal is to understand how visualization and interface design support situational awareness and efficient investigation.

Domain Area: ES Deployment

This domain covers the architectural and operational considerations involved in introducing Enterprise Security into an environment. Candidates should understand the dependencies, planning considerations, and general deployment approach required to support a stable and usable implementation. This includes awareness of how the solution fits into a broader Splunk architecture and how deployment choices affect manageability and performance.

Domain Area: Installation and Configuration

This area focuses on the foundational tasks required to install and configure Enterprise Security correctly. Candidates should understand how the product is prepared for use, how core settings are established, and how

configuration decisions influence later security content and workflows. The emphasis is on administrative understanding of setup and operational readiness.

Domain Area: Validating ES Data

A key part of Enterprise Security administration is ensuring that ingested data is usable, properly mapped, and suitable for security use cases. Candidates should understand how to confirm that data is arriving as expected, how normalization supports consistent analysis, and why data quality directly affects alerting, investigation, and reporting. This domain reflects the importance of trustworthy data in any SIEM environment.

Domain Area: Custom Add-ons

This domain addresses the role of add-ons in extending data support and improving compatibility with security data sources. Candidates should understand why custom add-ons may be required, how they contribute to field extraction and normalization, and how they support integration with unique or nonstandard technologies. The goal is to recognize how extensibility helps Enterprise Security adapt to real-world environments.

Domain Area: Tuning Correlation Searches

Candidates should understand how correlation searches are refined to improve relevance, reduce noise, and better reflect organizational risk. This includes the reasoning behind tuning detection logic, adjusting thresholds, improving context, and balancing sensitivity with operational practicality. Effective tuning is important because overly broad or poorly calibrated detections can reduce analyst efficiency and confidence.

Domain Area: Creating Correlation Searches

This area focuses on the design of new detection logic within Enterprise Security. Candidates should understand how correlation searches are used to identify suspicious activity patterns, transform search logic into operational detections, and generate actionable findings for security teams. The emphasis is on understanding how detection content is structured and how it supports meaningful security monitoring.

Domain Area: Lookups and Identity Management

Candidates should understand how contextual data improves security analysis. This includes the use of lookups to enrich events and the management of identity-related information to link activity to users, systems, or organizational roles. The domain highlights how context helps prioritize findings, improve investigations, and reduce ambiguity in security operations.

Domain Area: Threat Intelligence Framework

This domain covers the use of threat intelligence inside Enterprise Security to enhance detection and investigation. Candidates should understand how intelligence sources can be incorporated into workflows, how indicators provide additional context, and how threat intelligence contributes to more informed security decisions. The focus is on operational use and conceptual integration rather than on memorizing source-specific details.

Detailed Knowledge Explanation

SPLK-3001 ES Introduction

Splunk Enterprise Security (ES) is an advanced security application designed to run atop the standard Splunk platform. From a strategic architectural perspective, it transforms Splunk Core from a general-purpose machine data engine into a functional Security Information and Event Management (SIEM) system. For modern Security Operations Centers (SOCs), ES is the critical layer used to operationalize security intelligence, providing the correlation logic and workflow frameworks necessary to significantly improve Mean Time to Detect (MTTD) and Mean Time to Respond (MTTR).

1. What is Splunk Enterprise Security (ES)?

Splunk ES is a specialized SIEM application that provides security-specific analytics, dashboards, and incident management capabilities. While Splunk Core handles indexing and searching across any machine data, ES provides the logic required to identify, prioritize, and manage cybersecurity threats across the enterprise.

2. Why is Splunk ES Important?

As modern threats like phishing, advanced malware, and insider attacks become more sophisticated, organizations require more than basic log aggregation. Splunk ES provides real-time detection through correlation searches and ensures adherence to compliance standards. It provides a structured environment that allows security teams to move from reactive log searching to proactive, context-aware threat management.

3. Core Concept: Security Posture Dashboard

The Security Posture Dashboard serves as the command center for SOC analysts, offering a high-level view of organizational security health. It displays real-time insights via Notable Events—alerts generated when correlation searches meet specific risk conditions. By leveraging risk scoring, the dashboard enables analysts to immediately identify urgent issues, ensuring that critical threats are prioritized over routine system noise.

4. Core Concept: Data Models & CIM (Common Information Model)

The Common Information Model (CIM) is the set of standards that ensures field consistency across disparate data sources. Because firewalls, EDRs, and Windows logs use different naming conventions for the same attributes (e.g., `src_ip` vs `SourceAddress`), CIM maps these to standardized fields. This enables Data Models (e.g., Authentication, Network Traffic) to function. An architect can then run a single, high-performance query—such as `| datamodel Authentication search`—to identify failed logins across every vendor technology in the stack simultaneously.

5. Key Component: Notable Events

Notable Events represent conditions that stand out from the millions of raw logs collected. These are not just standard search results; they are high-fidelity alerts created by correlation searches that require human

intervention. Each Notable Event includes essential metadata, such as risk scores, event descriptions, and direct links to the underlying raw data for forensic investigation.

6. Key Component: Incident Review Dashboard

The Incident Review Dashboard is the primary workspace for the triage process. It provides the tools necessary for team coordination, allowing analysts to view, filter, and assign events. It supports the full incident lifecycle by allowing users to change event statuses (e.g., New, In Progress, Resolved) and document findings through notes and tags.

7. Key Component: Assets and Identities

Splunk ES uses lookup tables to map raw data to human-readable business context. This identity enrichment is a force multiplier for investigations. Identifying that a failed login involves the "CFO's laptop" as opposed to a "guest printer" changes the urgency of the response. Architects must understand the core Urgency calculation: **Urgency = Severity (from the search) + Priority (from the asset/identity record)**.

8. Modular Composition of Splunk ES

ES is a modular platform composed of specific Security Domain Modules: Access, Endpoint, Network, Identity, and Threat. Each module is independent and powered by its own data models and correlation logic. This independence allows for a scalable strategy where organizations can enable and tune only the modules relevant to their specific infrastructure and risk profile.

9. Relationship Between Splunk Apps, Add-ons, and ES

The ES data pipeline follows a strict three-tier architecture:

1. **Technology Add-ons (TAs):** Normalize raw logs via field extraction and CIM tagging.
2. **Data Models (CIM):** Standardize the structure of the parsed data.
3. **ES Modules:** Utilize these models for analytics and visualization.

Properly formatted input via TAs is the absolute prerequisite for ES functionality. Without CIM-compliant data, correlation searches will fail, and dashboards will remain unpopulated.

10. ES Introduction Practice Question

Q1: What is the primary function of Splunk Enterprise Security (ES) compared to Splunk Core?
A. To act as a Security Information and Event Management (SIEM) layer on top of Splunk Core
B. To enhance data ingestion throughput and reduce storage requirements
C. To replace traditional firewalls in modern IT environments
D. To serve as a replacement for Splunk Core in security use cases

Q2: What is the main purpose of the Security Posture Dashboard in Splunk ES?
A. To allow administrators to configure data ingestion pipelines
B. To monitor storage usage across all Splunk indexes

C. To display only network firewall activity
D. To provide real-time insight into the organization's security health and highlight Notable Events

Q3: What is the purpose of the Common Information Model (CIM) in Splunk ES?
A. To eliminate the need for raw data ingestion
B. To automatically enrich all incoming logs with geolocation data
C. To normalize and standardize data from multiple sources for consistent analysis
D. To create custom dashboards based on system logs

Q4: What makes a Notable Event "notable" in Splunk ES?
A. It always involves activity from domain administrators
B. It represents data collected from external threat feeds only
C. It is any event that occurs in the authentication data model
D. It meets specific criteria defined in correlation searches and requires human attention

Q5: Which of the following best describes the function of the Incident Review Dashboard?
A. A view for visualizing all asset-based lookup entries
B. A centralized workspace to triage and manage Notable Events
C. A place where correlation searches are configured
D. A dashboard that tracks system performance and index growth

Q6: In the context of Splunk ES, what is contextual enrichment?
A. Enhancing raw event data with meaningful asset and identity information
B. Encrypting sensitive fields within indexed events
C. Compressing logs to improve storage efficiency
D. Using machine learning to predict user behavior

Q7: Which component in Splunk ES uses lookup tables to relate IP addresses to assets?
A. Risk Scoring Engine
B. Assets and Identities Framework
C. Glass Tables Module
D. Event Timeline Visualizer

Q8: Which of the following is a benefit of using data models in Splunk ES?
A. They store raw log data in compressed binary format
B. They reduce the need to deploy Universal Forwarders
C. They allow users to run accelerated searches across normalized data sets
D. They automatically trigger real-time alerts for known attack patterns

Q9: A Notable Event in Splunk ES is typically generated by:
A. A manual tag applied by an analyst
B. Raw data ingestion without preprocessing
C. A correlation search that matches defined conditions
D. A dashboard widget configuration

Q10: What role does Splunk ES play in reducing MTTD and MTTR?
A. It integrates directly with firewalls to block IPs

B. It replaces the need for human analysis in security operations
C. It filters out low-priority logs before indexing
D. It provides dashboards, correlation searches, and alerts that accelerate detection and response

SPLK-3001 ES Deployment

The deployment of Splunk ES requires meticulous planning to ensure the architecture matches organizational data load and performance requirements.

1. What is ES Deployment?

ES deployment is the strategic installation of the application on Splunk Enterprise. It demands careful consideration of resource allocation and performance tuning to handle the intensive workloads generated by continuous correlation and data model summarization.

2. Deployment Architecture

- **Single-instance:** Suitable only for lab or PoC environments. It lacks the scalability for production data loads.
- **Distributed:** The enterprise standard, separating Search Heads, Indexers, and Forwarders to support parallel processing and high availability.

2.1 SHC (Search Head Cluster) Support

Search Head Clustering (SHC) is highly recommended for large security teams. It provides horizontal scaling and redundancy, ensuring that the SOC remains operational even if a single search head fails.

3. Requirements for Deploying ES

Deployment success relies on strictly following software and hardware prerequisites:

- **Compatibility:** ES versions must align with the base Splunk version (e.g., ES 7.0 requires Splunk Enterprise 9.1+).
- **Hardware:** ES is resource-heavy. Architects must allocate multiple high-performance cores, a minimum of 16–32 GB of RAM, and fast SSD storage to handle the high Disk I/O required for summary indexing and search workloads.

4. Best Practices for Deployment

For optimal performance, ES should run on dedicated Search Heads. Furthermore, architects must allocate separate indexes for security data (e.g., `notable`, `risk`, `datamodel_*`). This segregation improves search speed, simplifies data retention policies, and ensures better audit compliance.

5. Why Use Dedicated Indexes: Practical SPL Example

Searching dedicated indexes is exponentially faster than broad index searches. For example, a search focusing on security-specific data executes rapidly:

index=notable OR index=risk | stats count by source, sourcetype

This query ignores noisy, non-security data, providing analysts with immediate clarity while reducing the load on the indexer tier.

6. Data Model Acceleration and System Resource Impact

Data Model Acceleration (DMA) is critical for ES dashboards but places a heavy load on CPU and memory during summarization. Architects must stagger acceleration windows using cron scheduling and use the Monitoring Console (MC) to identify "resource-heavy data models."

7. ES Deployment Practice Question

Q1: What is the key difference between a single-instance and a distributed deployment of Splunk ES?
A. Distributed deployments separate components across multiple servers
B. Single-instance deployments support Search Head Clustering
C. Distributed deployments do not require indexers
D. Single-instance deployments are more scalable

Q2: Why is it recommended to use a dedicated search head for Splunk ES?
A. To simplify Splunk Core license management
B. To ensure compatibility with third-party firewall software
C. To isolate security operations and improve performance
D. To allow multiple Splunk apps to be installed simultaneously

Q3: What is the primary purpose of the Search Head Cluster (SHC) in an ES deployment?
A. To support active-passive failover of indexers
B. To eliminate the need for data model acceleration
C. To reduce forwarder management complexity
D. To provide high availability and horizontal scaling for search heads

Q4: What must be reviewed before installing a new version of Splunk ES?
A. Risk scoring policy
B. Daily search quota
C. Compatibility matrix with the current Splunk Core version
D. Index volume threshold

Q5: What hardware resource is especially critical for Splunk ES due to data model acceleration and summary indexing?
A. Disk I/O performance
B. GPU performance

C. Network bandwidth
D. CPU threading efficiency

Q6: Which of the following indexes is typically used to store correlation search results?
A. _internal
B. audit
C. summary
D. notable

Q7: In a distributed Splunk ES deployment, what is the role of indexers?
A. To store and process incoming event data
B. To schedule risk rule execution
C. To collect and forward syslog data
D. To manage user permissions and interface preferences

Q8: What is one reason to use separate indexes for security data like `notable` and `risk`?
A. To reduce the need for data model acceleration
B. To improve data segmentation and compliance tracking
C. To apply different license policies to each index
D. To enable log forwarding to external SIEMs

Q9: What is the minimum recommended RAM for a Splunk ES production search head?
A. 4 GB
B. 16 GB
C. 8 GB
D. 64 GB

Q10: Why is it not recommended to install Splunk ES on a shared search head?
A. It disables the Security Posture Dashboard
B. It can reduce search performance due to resource competition
C. It causes automatic license violations
D. It prevents the use of search head clustering

SPLK-3001 Installation and Configuration

Installation and configuration represent a multi-step process involving dependency verification and post-install validation to ensure environment stability.

1. What is Installation and Configuration in Splunk ES?

This phase ensures ES integrates efficiently with the Splunk setup. Proper configuration prevents ingestion bottlenecks and ensures that the analytics engine is functional.

2. Installation Methods

- **Splunkbase GUI:** Preferred for smaller, standalone instances.
- **CLI:** The preferred method for automated or distributed deployments. Example: `splunk install app /tmp/SplunkEnterpriseSecurity.spl -auth admin:password`

3. Dependency Checks

ES relies on several supporting apps (SAs) and domain add-ons (DAs). Critical dependencies that must be present include **SA-CIM**, **DA-ESS-Content**, **SA-ThreatIntelligence**, **SA-Utils**, and **SA-Eventgen**. Missing any of these will result in broken backend operations and missing features.

4. Post-installation Checks

After restarting Splunk, administrators must run health checks via the Monitoring Console (MC). The exact path is `Splunk Enterprise Security > Health Check`. Key areas to monitor include KV Store status, search scheduler queue health, and indexing throughput.

5. Configuration: Data Model Acceleration (DMA)

Admins must enable DMA for required models (e.g., Authentication, Intrusion Detection) and set summary ranges. DMA is mandatory for real-time correlation searches but must be monitored for disk storage consumption.

6. Configuration: User Roles and Permissions

ES uses Role-Based Access Control (RBAC) to protect the environment:

- **ess_admin:** Full control over detection logic and risk rules.
- **ess_user:** Limited to viewing dashboards and performing investigations.

7. Additional Tasks: Assets and Identities Mapping

To enable context-aware investigations, admins must upload and map `assets.csv` and `identities.csv`, ensuring fields like IP, MAC, user, and priority are correctly populated.

8. Additional Tasks: Notable Event Aggregation Policies

Aggregation policies group related events to reduce noise. **Example Aggregation Rule:**

- **Group By:** `user`, `src_ip`
- **Time Window:** 5 minutes
- **Outcome:** 10 failed login attempts from the same source result in one logical alert for the analyst.

9. Installation and Configuration Practice Question

Q1: Which method is appropriate for installing Splunk Enterprise Security in a distributed environment?
A. Using the command line with the install app command
B. Installing through the Windows Services console
C. Uploading via the Splunk Add-on Builder
D. Downloading from Splunkbase and double-clicking the .spl file

Q2: What is the function of the `SA-CIM` component in Splunk ES?
A. Acts as the default dashboard template
B. Enables search head clustering
C. Provides default indexes for ES content
D. Supports the Common Information Model (CIM)

Q3: What should an administrator verify immediately after installing Splunk ES?
A. That no raw data is being indexed
B. That all internal Splunk licenses are deactivated
C. That every indexer has been restarted
D. That the navigation menu and dashboards load properly

Q4: What is the role of Data Model Acceleration (DMA) in Splunk ES?
A. Automatically deletes old notable events
B. Compresses log files for cheaper storage
C. Increases indexer cluster replication factor
D. Improves dashboard speed and enables real-time correlation

Q5: Which user role in Splunk ES has full administrative access, including to correlation searches?
A. ess_user
B. ess_admin
C. admin_readonly
D. ess_operator

Q6: Why is enabling acceleration for relevant data models critical in Splunk ES?
A. It allows for the deletion of raw event data
B. It reduces the number of indexers required
C. It removes the need for forwarders
D. It enables real-time dashboards and correlation searches

Q7: What is the purpose of uploading `assets.csv` and `identities.csv` in ES configuration?
A. To enrich events with contextual information
B. To activate archiving policies
C. To enhance licensing and quota control
D. To tune correlation search thresholds

Q8: How does configuring Notable Event Aggregation Policies benefit a security team?
A. By clustering indexers for alert sharing
B. By grouping related alerts into incidents and reducing alert noise

C. By replacing all raw event storage
D. By forwarding all alerts to third-party SIEMs

Q9: What is a key reason to configure SMTP settings after installing Splunk ES?
A. To notify security teams when critical alerts are triggered
B. To enforce license usage quotas
C. To automatically restart indexers
D. To send enriched identity data to Splunk Core

Q10: Which of the following is a best practice when assigning user roles in Splunk ES?
A. Assign only the `admin` role to every analyst
B. Disable all user roles except `ess_admin`
C. Use role inheritance to define custom roles
D. Allow all users to modify correlation search logic

SPLK-3001 Custom Add-ons

Custom Technology Add-ons (TAs) are essential for normalizing non-standard vendor data to ensure CIM compliance.

1. What Are Custom Add-ons in Splunk ES?

Custom TAs extract fields and apply tags from raw data when a vendor does not provide an official TA. They act as the parser that prepares data for ES analytics.

2. Key Tasks in Building Custom Add-ons

Using the Splunk Add-on Builder, developers create data inputs (Syslog, REST API), define field extractions using regex, and assign event types and tags.

3. CIM Compliance

Field naming must follow CIM standards exactly: use `src` (not `src_ip`) and `user` (not `username`). Required tags (e.g., `tag=authentication`) must be applied to identify the event function.

4. Packaging and Deployment

TAs must be lightweight and contain no UI elements. They are deployed to Universal Forwarders for collection or Heavy Forwarders/Search Heads for parsing.

5. Testing and Validation

Use the CIM Validation Dashboard and Data Model Audit Tool to identify gaps in tagging or extraction before pushing to production.

6. Sample Directory Structure for a Technology Add-on (TA)

Standardized structures ensure compatibility with deployment tools:

- `bin/`: Input/enrichment scripts.
- `default/`: Core configurations (props, transforms).
- `local/`: Environment-specific overrides (do not package).
- `metadata/`: Permissions (ensure `default.meta` allows global access).

7. Manual Review of Auto-Generated Config Files

The Add-on Builder is a starting point, but manual auditing of `props.conf` and `transforms.conf` is **mandatory**. Specifically, you must verify `TIME_FORMAT` and `MAX_TIMESTAMP_LOOKAHEAD` settings, as incorrect timestamping is the primary cause of data fragmentation.

8. Custom Add-ons Practice Question

Q1: Why would you need to create a custom Technology Add-on (TA) in Splunk ES?
A. To normalize data and apply CIM tags when no official TA exists
B. To visualize dashboards from unsupported vendors
C. To bypass role-based access control in Splunk
D. To enable data forwarding from third-party log management tools

Q2: Which tool is primarily used to build custom add-ons for Splunk?
A. Event Type Aggregator
B. Risk Scoring Manager
C. Add-on Builder
D. CIM Compliance Validator

Q3: Which of the following field names is CIM-compliant?
A. `source_ip`
B. `src`
C. `username`
D. `ip_src`

Q4: Why are tags like `tag=authentication` important in a custom TA?
A. They are used to throttle alert generation
B. They define index retention policies
C. They control user role access to dashboards
D. They help map events to CIM data models

Q5: Where should a custom TA be deployed if it performs field extractions using props and transforms?
A. On heavy forwarders or search heads
B. Only on the license master
C. Only on Universal Forwarders
D. On index replication peers

Q6: Which type of input can be configured in the Add-on Builder?
A. Modular inputs for REST API polling
B. JavaScript SDK triggers
C. Real-time distributed database querying
D. Dashboard panels

Q7: Which of the following would violate best practices when packaging a TA?
A. Using field aliases for CIM compliance
B. Adding event types and tags
C. Including props.conf and transforms.conf
D. Embedding dashboards and saved searches

Q8: What does the CIM Validation Dashboard in the Add-on Builder help you confirm?
A. Whether your field extractions and tags are populating expected data models
B. Whether your TA has been uploaded to Splunkbase
C. Whether the data model summary index is optimized
D. Whether your add-on includes at least three tags per event

Q9: What is the purpose of using field aliases in a custom TA?
A. To reduce the number of sourcetypes
B. To allow non-CIM fields to be used in CIM-based searches
C. To increase licensing efficiency
D. To prevent duplicate tags from triggering correlation searches

Q10: What should you review in the Data Model Audit Dashboard after deploying a new custom TA?
A. Whether events from the TA are populating the appropriate data models
B. Number of tags defined in default.meta
C. Forwarder CPU utilization
D. License usage by source type

SPLK-3001 Lookups and Identity Management

Lookups provide the organizational context—the who, what, and where—required to transform raw logs into actionable intelligence.

1. What Are Lookups in Splunk ES?

Lookups are reference tables used to map raw values (like IPs) to business context (like a department or asset owner).

2. Types of Lookups

- **CSV File Lookups:** Static mappings.
- **KV Store Lookups:** Dynamic, scalable, and editable via the UI/API.
- **External Lookups:** Python scripts querying external APIs (e.g., VirusTotal) in real-time.

3. Identity Management in Splunk ES

Identity management utilizes `assets.csv` and `identities.csv`. The **Priority** field (Low to Critical) in these files is vital, as it directly scales the risk score of alerts involving those entities.

4. Automatic Lookups: props.conf + transforms.conf

Automatic lookups enrich data at search time. `transforms.conf` defines the logic, while `props.conf` determines which sourcetypes receive the enrichment, eliminating the need for manual `| lookup` commands in analyst queries.

5. Handling Lookup Failures (Fallback Behavior)

If a match is not found, Splunk returns a null value. This behavior ensures query robustness, and analysts can handle gaps using `coalesce` or `fillnull` in SPL.

6. Using the Splunk Web GUI: Lookup Editor for KV Store Lookups

The Lookup Editor app allows admins to manage KV Store collections directly from the web interface, which is the preferred method for updating identity or asset records without CLI access.

7. Lookups and Identity Management Practice Question

Q1: What is the primary purpose of using lookups in Splunk ES?
A. To replicate dashboards between environments
B. To reduce the number of correlation searches in use
C. To enrich event data with additional organizational or business context
D. To accelerate field extraction on indexers

Q2: Which file is typically used in Splunk ES to map usernames or emails to departments and access tiers?
A. `assets.csv`
B. `priorities.conf`
C. `accounts.csv`
D. `identities.csv`

Q3: What type of lookup is best suited for tracking a dynamic inventory of assets that changes frequently?
A. KV Store lookup
B. External lookup
C. CSV lookup
D. Manual lookup

Q4: What is a typical use case for an external lookup in Splunk ES?
A. Manually mapping usernames to email addresses
B. Automatically updating the `notable` index
C. Viewing field extractions in raw log view
D. Calling an external API to enrich IP data with live threat intelligence

Q5: In which directory are the core identity lookup files (`assets.csv`, `identities.csv`) typically located?
A. `$SPLUNK_HOME/bin/system/lookups`
B. `$SPLUNK_HOME/etc/apps/SA-IdentityManagement/lookups`
C. `$SPLUNK_HOME/etc/system/default/lookups`
D. `$SPLUNK_HOME/etc/apps/SA-CIM/lookups`

Q6: Which field in an identity or asset record influences correlation search risk scoring?
A. `type`
B. `priority`
C. `source`
D. `tag`

Q7: What functionality does the Asset and Identity Center provide in Splunk ES?
A. Reviewing and validating loaded identity and asset records
B. Scheduling updates to correlation search macros
C. Creating notable events directly from lookups
D. Editing raw data model configurations

Q8: Why might a lookup-enhanced correlation search be more effective than a raw event search?
A. It removes the need for indexing rules
B. It ensures searches only use event types
C. It can enrich alerts with asset roles and business impact context
D. It suppresses alerts for test systems

Q9: What is a typical format of a CSV lookup used in Splunk?
A. `field1,field2\nvalue1,value2`
B. `{ "key": "value" }`
C. `lookup=value1|value2|value3`
D. `field1:value1; field2:value2`

Q10: Which lookup type would best support real-time editing by analysts without restarting Splunk?
A. CSV lookup
B. KV Store lookup

C. External lookup
D. Static JSON file

SPLK-3001 Creating Correlation Searches

Correlation searches are the primary detection mechanism, identifying suspicious patterns across multiple data sources in real-time.

1. What Are Correlation Searches?

These are scheduled SPL queries that generate Notable Events or trigger Adaptive Response Actions (ARAs).

2. Core Components: Search Logic (SPL Development)

Architects must use efficient SPL commands such as `tstats`, `eval`, `stats`, and `datamodel`. Logic should be modular and specific to avoid resource drain.

3. Core Components: Correlation Search Editor

The editor allows for the configuration of Severity, Urgency, and Category. It supports dynamic tokens (e.g., `$user$`, `$src$`) in titles and descriptions to provide immediate context in Incident Review.

4. Core Components: Integration with Other Data

Effective searches leverage the Threat Intelligence Framework and Identity context to improve accuracy and relevance.

5. Event Type Registration and CIM Alignment

Events must be registered with appropriate event types (e.g., `failed_logins`) and tags (e.g., `tag=authentication`). This is critical for ensuring data populates the Security Posture Dashboard.

6. Behind the Scenes: adaptive_response_actions.conf

ARAs are defined in `adaptive_response_actions.conf`. To allow global access for custom actions, architects must ensure `default.meta` permissions are set correctly.

7. Clarifying: Schedule vs. Time Range vs. Throttle

Understanding the interaction between timing and suppression is critical for a stable SIEM.

Component	Description
Schedule (cron)	Defines when the search runs (e.g., `*/5 * * * *` for every 5 mins).
Earliest/Latest	Defines the time window of data examined (e.g., `-5m` to `now`).
Throttle	Prevents repeated alerts for the same entity (e.g., `user`) for a set period.

8. Creating Correlation Searches Practice Question

Q1: What is a key purpose of the Correlation Search Editor in Splunk ES?
A. To route data from forwarders into risk indexes
B. To define event categories, severities, and response actions
C. To manage data model acceleration and summaries
D. To configure indexer clustering and failover

Q2: Which of the following best describes a common use of tokens like `$user$` or `$src$` in a correlation search?
A. Dynamically inserting field values into titles or descriptions
B. Encrypting notable event descriptions
C. Linking the search to CIM model acceleration
D. Generating lookup tables at runtime

Q3: Which SPL structure is best suited for efficiently searching a large volume of authentication failure logs in a correlation search?
A. `| inputlookup` across identity models
B. `| transaction` with default parameters
C. `| eventstats` on all fields without summary
D. `| datamodel Authentication Authentication search` with filters

Q4: Why should correlation searches be tested in the Search & Reporting app before being deployed?
A. To create permanent lookup references
B. To automatically trigger the notable event workflow
C. To store SPL macros centrally
D. To validate logic, field extractions, and result volume

Q5: What is the purpose of defining urgency and severity when configuring a correlation search?
A. To calculate data retention time
B. To reduce the volume of indexed logs
C. To prioritize events in dashboards and assign risk impact
D. To track licensing across ES modules

Q6: What does the Content Management dashboard in Splunk ES provide visibility into?
A. Authentication failures across all indexes
B. Threat feeds and adaptive response scripts
C. A list of correlation searches and their execution stats
D. Field aliases and lookup sources only

Q7: What is the role of the `lookup` command in correlation searches?
A. To normalize CIM fields using indexed metadata
B. To manually group transactions based on event types
C. To enrich events using external datasets or threat lists
D. To create auto-scheduled search macros

Q8: Which of the following represents a best practice when building a correlation search?
A. Leverage tstats or datamodels where applicable
B. Use full subsearch joins on multiple sourcetypes
C. Run the search hourly using real-time mode
D. Skip eventtype tagging to improve speed

Q9: In the Correlation Search Editor, how can you integrate an external ticketing workflow when a Notable Event is triggered?
A. Use an Adaptive Response Action to call a webhook or script
B. Schedule an inputlookup refresh
C. Manually tag the event in the data model
D. Assign a risk category directly to an asset

Q10: Which type of integration can be used in a correlation search to determine whether an IP address belongs to a known malicious actor?
A. Applying event sampling on indexed data
B. Referencing a threat intelligence lookup file
C. Using `rex` to extract severity from logs
D. Clustering alerts from multiple dashboards

SPLK-3001 Tuning Correlation Searches

Tuning is an ongoing operational necessity to maintain system performance and prevent analyst alert fatigue.

1. Why Is Tuning Crucial?

Broad searches generate excessive noise and consume CPU/RAM. Tuning ensures the SOC focuses only on high-fidelity, actionable threats.

2. Tuning Techniques: Modify Search Criteria

Narrow time ranges and add specific filters, such as focusing searches only on production systems or high-privilege accounts.

3. Tuning Techniques: Adjust Scheduling

High-cost searches should be spread using cron expressions to avoid overloading the scheduler. Throttling is essential; for example, setting a 30-minute throttle period for a `user` field ensures one alert for a brute-force attempt rather than fifty.

4. Tuning Techniques: Suppression Rules

Suppression rules hide expected activity (e.g., vulnerability scans) from the Incident Review dashboard, keeping it focused on actual incidents.

5. Risk Scoring Optimization

Moving to Risk-Based Alerting (RBA) allows for the aggregation of multiple low-risk events. A Notable Event is only triggered once a risk threshold is crossed, significantly improving the signal-to-noise ratio.

6. Tuning Correlation Searches Practice Question

Q1: What is the primary purpose of tuning correlation searches in Splunk ES?
A. To reduce alert fatigue and optimize system resource usage
B. To convert SPL to JSON for performance reasons
C. To disable low-urgency dashboards
D. To enhance raw log ingestion rates

Q2: Which method best helps to avoid generating repetitive alerts for the same activity within a short time frame?
A. Using scheduled search acceleration
B. Adjusting the search priority level
C. Setting throttle settings
D. Enabling summary indexing

Q3: A correlation search is configured to run every 5 minutes but searches the last 24 hours of data. What is the most likely consequence?
A. Delayed indexing of events
B. Dashboard load failure
C. High system resource consumption
D. Failure to tag risk objects

Q4: Which scenario would most benefit from suppression rules in Splunk ES?
A. A known vulnerability scan running weekly
B. A newly installed app's initial data ingestion
C. A search head restart incident
D. A sudden increase in license usage

Q5: What does tuning the correlation search schedule primarily help you control?
A. Indexer retention policies
B. Frequency and system load of detection logic
C. Timing of field extractions
D. User access to notable events

Q6: What is the purpose of assigning risk scores instead of immediately generating notable events?
A. To archive lower-priority events directly
B. To accumulate contextual risk for prioritized alerts
C. To reduce the need for data model acceleration
D. To prevent searches from being stored in audit logs

Q7: Which of the following is a recommended technique for improving correlation search performance?
A. Use wildcards in every SPL condition
B. Expand the search window as much as possible
C. Increase the frequency to every 1 minute
D. Add filters for severity, source type, or critical assets

Q8: How does customizing risk thresholds for different users or systems improve correlation search effectiveness?
A. It enables alert encryption
B. It ensures risk is assessed based on asset criticality
C. It helps generate alerts more uniformly
D. It guarantees correlation search acceleration

Q9: Which configuration change can help spread correlation search load evenly over time?
A. Modify cron scheduling
B. Throttle on index names
C. Reboot the search head cluster
D. Update the event type tags

Q10: What is the most likely result of failing to tune correlation searches in a production ES environment?
A. Analysts may miss critical threats due to alert overload
B. Risk scores will reset every 15 minutes
C. Dashboards will only show historical data
D. Event types will stop updating automatically

SPLK-3001 Security Intelligence

Security Intelligence shifts the focus from simple alerting to context-aware risk management.

1. Core Concept 1: Risk-Based Alerting (RBA)

RBA accumulates risk scores for entities over time. This focuses the SOC on overall behavioral patterns rather than individual anomalies.

2. Core Concept 2: Risk Rules

Risk Rules assign numeric scores to entities and store them in the `risk` index. They do not alert directly but build the cumulative score used for RBA.

3. Core Concept 3: Assets and Identities Framework

This framework provides the context needed to treat a "Domain Controller" differently than a "Guest Device," applying higher risk weights to critical organizational components.

4. How Risk Rules Are Configured and Stored

Risk Rules are managed via the Content Management UI. Scored risk events are stored in the internal `risk` index, containing fields like `risk_object`, `risk_score`, and `rule_name`.

5. SPL Example: Identify High-Risk Users

Architects use `tstats` for high-performance querying of the Risk data model:

| tstats sum(All_Risk.risk_score) as total_risk from datamodel=Risk.All_Risk by All_Risk.user | where total_risk > 50

6. Security Intelligence Practice Question

Q1: What is the primary goal of Security Intelligence in Splunk ES?
A. To limit access to the Incident Review dashboard
B. To reduce alert noise and help analysts focus on meaningful threats
C. To assign severity labels to correlation searches
D. To automatically enrich every log with geographic information

Q2: What is a key feature of Risk-Based Alerting (RBA) in Splunk ES?
A. It accumulates risk scores for entities and only alerts when thresholds are met
B. It triggers a notable event every time an event occurs
C. It suppresses all alerts unless triggered by an administrator
D. It replaces correlation searches with prebuilt dashboards

Q3: What role does the Assets and Identities Framework play in Security Intelligence?
A. It deletes non-CIM-compliant data before enrichment
B. It indexes raw data into the CIM model
C. It provides context by linking technical identifiers to business-relevant information
D. It stores search history for high-value users

Q4: Which of the following best describes a Risk Rule in Splunk ES?
A. A lookup file that lists all valid user accounts
B. An index that stores threat intelligence feeds
C. A dashboard component that shows system health metrics
D. A type of correlation search that adds risk scores instead of triggering alerts

Q5: Which behavior is most likely to have a higher risk score in an RBA model?
A. A logout from a regular user account
B. A port scan on a domain controller
C. A failed login from a guest printer
D. A DNS request from a kiosk machine

Q6: What does Splunk ES do when multiple Risk Rules apply to a single entity?
A. Overwrites the risk score with the most recent rule
B. Automatically disables the least severe rule
C. Sums the risk scores to determine overall threat level
D. Suppresses all alerts until all rules are cleared

Q7: Why is contextual enrichment important in Security Intelligence?
A. It compresses data models to improve performance
B. It helps translate event logs into high-urgency events based on business value
C. It identifies irrelevant log sources
D. It allows correlation searches to bypass raw data

Q8: Which of the following is most likely stored in the `risk` index in Splunk ES?
A. Asset information including MAC addresses
B. Aggregated risk scores assigned by risk rules
C. Raw notable event logs
D. Data model acceleration status

Q9: What is a main benefit of using Risk Rules rather than standard correlation searches?
A. They allow gradual risk accumulation instead of immediate alerts
B. They generate more alerts to ensure full coverage
C. They bypass CIM compliance checks for faster detection
D. They can be configured directly from the command line only

Q10: Which statement best describes how Security Intelligence reduces alert fatigue?
A. It correlates events, scores risk, and only escalates critical threats
B. It requires analyst approval before logging any event
C. It tags all events with “low priority” by default
D. It filters out alerts based on index size

SPLK-3001 Threat Intelligence Framework

The Threat Intelligence Framework (TIF) enables the ingestion and operationalization of external Indicators of Compromise (IOCs).

1. What is the Threat Intelligence Framework (TIF)?

TIF matches malicious IPs, domains, and hashes against internal logs to provide actionable alerts.

2. Supported Threat Sources

TIF supports TAXII feeds (STIX), custom CSV/JSON files, and commercial platforms (e.g., Recorded Future, Anomali).

3. Integration with ES Detection Logic

IOCs are stored in KV Store lookups and used in correlation searches. Matches can trigger Notable Events or, more scalably, increment risk scores via RBA.

4. Framework Components: Manager Dashboard and KV Store

The Threat Intelligence Manager monitors feed health. The KV Store supports high-speed lookups and allows for "aging" policies.

5. Key Features: Threat Match Search and Enrichment

Matched events are enriched with metadata, including the threat source, type, and confidence level.

6. IOC Aging Strategy and Validation

Freshness is maintained through per-type aging policies:

- **IP Address:** 7 days
- **Domain Name:** 14 days
- **File Hash:** 30 days

7. Threat Intelligence Framework Practice Question

Q1: What is the primary function of the Threat Intelligence Framework (TIF) in Splunk ES?
A. To monitor index usage and detect retention violations
B. To configure authentication methods and SSO policies
C. To collect and match threat indicators (IOCs) against enterprise log data
D. To manage Splunk license allocation across search heads

Q2: Which of the following feed types is directly supported by the TIF using a structured threat exchange protocol?

A. TAXII feeds containing STIX-formatted IOCs
B. syslog-formatted logs
C. JSON uploads from Splunkbase
D. CSV lookups from search head peers

Q3: Which of the following lookup tables would most likely contain IOC data related to malicious file hashes?
A. threat_intel_by_email
B. risk_by_asset
C. threat_intel_by_file_hash
D. threat_intel_by_ip

Q4: In the Threat Intelligence Framework, what is the primary function of KV Store tables?
A. To override CIM field mappings at runtime
B. To allow indexers to ingest threat data directly
C. To store, expire, and enrich IOC records dynamically
D. To manage event correlation search acceleration

Q5: What is the purpose of the Threat Intelligence Manager dashboard?
A. To configure lookup table permissions
B. To audit login activity of ES admins
C. To display correlation search results
D. To validate the status of threat feeds and view recent IOC ingestion activity

Q6: What field is commonly used to evaluate the reliability of an IOC in TIF?
A. event_type
B. priority
C. confidence
D. action

Q7: Which of the following use cases best fits a threat match search in Splunk ES?
A. Detecting field extraction misconfigurations
B. Finding current log activity that matches known IOCs
C. Aggregating license usage data
D. Identifying search head cluster failures

Q8: How does TIF improve the quality of notable events generated by correlation searches?
A. It enriches events with IOC metadata like source, type, and threat confidence
B. It bypasses the Common Information Model (CIM) for faster results
C. It replaces adaptive response actions with static alert templates
D. It allows low-confidence IOCs to be ignored automatically

Q9: Which of the following formats is valid for manually uploading threat intel to Splunk ES?
A. CSV or JSON files containing IPs, domains, or file hashes
B. Proprietary binary logs
C. Custom macros embedded in dashboards
D. Raw packet captures

Q10: Which field in a threat intel record typically determines how long it will remain active in TIF?
A. ttl (time to live)
B. first_seen
C. correlation_id
D. event_time

SPLK-3001 Monitoring and Investigation

Monitoring and investigation are the daily operational rhythm of the SOC, where alerts are triaged and resolved.

1. Feature 1: Incident Review

The central hub for alert management, filtering, and assignment.

2. Feature 2: Event Timeline

A visual narrative tool that provides a chronological view of events across multiple sources, essential for root-cause analysis.

3. Feature 3: Search & Drilldown Capabilities

Allows analysts to pivot from a summary Notable Event directly into the raw log data to uncover the full scope of an attack.

4. Use Case: Failed Login Analysis

Analysts triage events by checking risk scores, using the timeline to see surrounding activity, and drilling into raw logs to identify brute-force patterns or lateral movement.

5. Understanding Urgency Levels

Urgency is the product of severity and asset priority. This ensures a High Severity alert on a Critical system is addressed first.

6. Adaptive Response Actions (ARAs) in Workflow

ARAs automate responses—such as blocking an IP or creating a ticket—to reduce MTTR.

7. Monitoring and Investigation Practice Question

Q1: What is the primary purpose of the Incident Review dashboard in Splunk ES?
A. To manage and coordinate the review of Notable Events

B. To configure data ingestion pipelines and source types
C. To track Splunk license usage across all indexes
D. To create CIM-compliant dashboards for operational monitoring

Q2: Which of the following data elements is typically included in a Notable Event in Splunk ES?
A. Storage tier classification and index replication factor
B. Event name, time, risk score, and links to related users or systems
C. Configuration of data model acceleration
D. List of raw logs from all data sources

Q3: What is the function of the Event Timeline feature in Splunk ES?
A. To store raw logs associated with Notable Events
B. To summarize license usage over time
C. To create adaptive response rules for correlation searches
D. To visualize the temporal relationship between events

Q4: Which capability allows analysts to directly pivot into deeper investigations within Splunk ES?
A. Glass Tables
B. Risk-based alerting
C. Drilldown and custom SPL search
D. Alert suppression filters

Q5: In a use case involving repeated failed logins, what would be a logical next step for an analyst after identifying a Notable Event?
A. Archive the raw logs for long-term storage
B. Restart the Splunk indexers to ensure data freshness
C. Investigate the source IP, review login attempts, and escalate if needed
D. Disable correlation search to reduce alert noise

Q6: How can the use of tags and status changes within Incident Review improve investigation workflow?
A. They automatically enrich all events with threat intelligence
B. They determine whether an event can be routed to Splunk Core
C. They increase the frequency of correlation search execution
D. They help classify and track the progress of events

Q7: What best describes the purpose of filtering options in the Incident Review dashboard?
A. To help analysts prioritize events based on time, urgency, or ownership
B. To restrict correlation search execution to working hours
C. To limit the storage size of Notable Events
D. To select only events tied to specific source types

Q8: When might an analyst choose to close a Notable Event without escalation?
A. When the indexer queue length exceeds 90%
B. When the source IP is a trusted internal host and the activity is expected
C. When no correlation search has fired for that user within 30 minutes
D. When the data model used is no longer accelerated

Q9: What is a key benefit of using Event Timeline during threat investigations?
A. It simplifies search string syntax by visualizing SPL results
B. It helps determine the order and connection between events
C. It configures correlation search thresholds more easily
D. It displays Splunk system logs side by side with event data

Q10: Which of the following best demonstrates a precision investigation using Splunk ES?
A. Pivoting from a Notable Event to a raw log search using field-based drilldown
B. Filtering notable events using correlation search priority levels
C. Reviewing notable events in Splunk Light
D. Monitoring disk I/O performance of Splunk indexers during alert generation

SPLK-3001 Forensics, Glass Tables, and Navigation Control

Advanced tools provide deep investigation for analysts and strategic visualization for leadership.

1. Forensics: Raw Event Access and Search Reconstruction

Analysts move beyond alert summaries to raw logs to build a complete sequence of an attacker’s actions.

2. Glass Tables: Strategic KPI Visualization

Real-time, interactive dashboards that visualize high-level security posture and KPIs for executives and SOC managers.

3. Navigation Control: Role-Based Experience

Admins customize the UI based on the user's role to enforce workflow discipline.

Role Name	Default Landing Page	Accessible Features
`es_analyst`	Incident Review	Drilldown, Timeline, Investigation Tools
`es_manager`	Security Posture Dashboard	Glass Tables, KPI Dashboards, CS Configuration

4. Practical SPL for Timeline-Based Forensics

Stitch together multi-source logs using a unified chronological view:

index=windows OR index=paloalto | table _time, user, src, dest, action | sort _time

5. Forensics, Glass Tables, and Navigation Control Practice Question

Q1: What is the primary purpose of forensics within Splunk ES?
A. To schedule correlation searches at specific times
B. To export dashboards to external SIEMs
C. To investigate incidents using raw logs and event reconstruction
D. To manage Splunk licenses and user roles

Q2: Which feature allows analysts to view the raw data behind a Notable Event?
A. Glass Tables
B. Risk-Based Alerting
C. Incident Review drilldown
D. Timeline Correlation

Q3: What can the visual timeline view in Splunk ES help an analyst do?
A. See how events from multiple systems occurred in chronological order
B. Configure new data inputs
C. Review license consumption
D. Correlate unrelated dashboards

Q4: What type of user benefits most from Glass Tables?
A. Splunk system developers writing new add-ons
B. Executives and SOC managers looking for high-level security overviews
C. IT hardware maintenance technicians
D. Database administrators

Q5: Which of the following best describes a use case for Glass Tables?
A. Visually tracking threats to critical business units
B. Editing correlation search logic
C. Analyzing log events one by one
D. Tuning indexer performance

Q6: In Splunk ES, what is the role of Navigation Control?
A. To customize the user interface based on user roles
B. To map CIM fields across data models
C. To define search head clustering roles
D. To adjust licensing thresholds during peak hours

Q7: How can Navigation Control help new analysts?
A. By filtering out low-risk logs automatically
B. By guiding them through structured workflows and simplified interfaces
C. By limiting data model availability during search
D. By allowing full administrative access on login

Q8: When reviewing a potential threat, what does “search reconstruction” refer to in Splunk ES?
A. Rebuilding an attacker’s sequence of actions using log data
B. Exporting Splunk dashboards to external SIEMs
C. Visualizing traffic patterns with IP geolocation
D. Automatically re-indexing data for timeline accuracy

Q9: Which of the following is a benefit of using timeline views during forensic investigation?
A. They replace the Incident Review dashboard for alert triage
B. They prioritize search results by storage usage
C. They eliminate the need for correlation searches
D. They help analysts trace events across systems in time order

Q10: How can Glass Tables improve response time during an active threat?
A. By providing real-time visual indicators of impacted assets
B. By auto-deploying firewalls to compromised devices
C. By hiding low-severity events from analysts
D. By automatically disabling correlation searches

SPLK-3001 Validating ES Data

Validation is the final critical step to ensure CIM compliance and SIEM reliability.

1. What is “Validating ES Data”?

Confirming that data is parsed, tagged, and mapped correctly to CIM data models.

2. Validation Tools: Data Model Audit and datamodel SPL

Use the Data Model Audit dashboard for coverage percentages. Use the `| datamodel` command to manually verify model population:

| datamodel Authentication Authentication_Failed search

3. Common Data Validation Issues

Blank panels often stem from incorrect field naming (e.g., `username` vs `user`), missing required tags (e.g., `tag=malware`), or index permission issues.

4. Using Search Inspector for Manual Field Validation

If a dashboard is blank, use the Search Inspector to diagnose extraction issues or verify if fields are missing at search time.

5. Best Practices for ES Data Validation

Integrity is maintained by using sample events for manual validation and creating test correlation searches. This ensures the entire ES lifecycle remains functional and secure.

6. Validating ES Data Practice Question

Q1: What is the purpose of the Data Model Audit Dashboard in Splunk ES?
A. To configure summary indexing for dashboards
B. To monitor forwarder health status
C. To automatically resolve event-type conflicts
D. To validate data coverage, tagging, and CIM compliance

Q2: Which SPL command helps validate that data is being mapped to the correct data model?
A. `| datamodel <model> <object> search`
B. `| metadata type=datamodel`
C. `| inputlookup datamodel_status`
D. `| tstats count by sourcetype`

Q3: Which of the following issues would most likely prevent failed login events from appearing in the Authentication data model?
A. The user role is not `ess_user`
B. The user is not in the identities.csv file
C. The data lacks the tag `authentication`
D. The events are stored in an accelerated index

Q4: A firewall log includes the field `src_ip` instead of `src`. What is the impact on Splunk ES?
A. The field will be ignored by dashboards expecting `src`
B. The value will be duplicated and slow down indexing
C. Correlation searches will run faster
D. The log will be indexed twice by accident

Q5: What is one way to manually verify whether a raw event is CIM-compliant?
A. Run the SPL query with `eventstats`
B. Restart the Splunk search head and review memory usage
C. Use the Search Inspector to validate field and tag presence
D. Check if it appears in internal logs

Q6: What is the risk if a correlation search is configured correctly but returns no results?
A. The dashboard is cached and needs a manual refresh
B. Data may be stored in an index not visible to the ES app

C. The Splunk license has expired
D. The search head cluster is not synchronized

Q7: Why is it important to check the presence of event types and tags in data onboarding?
A. Because ES dashboards are filtered by role
B. Because they determine whether data maps into a data model
C. Because they allow searches to bypass acceleration
D. Because tags speed up indexing throughput

Q8: What is the function of the `identities.csv` and `assets.csv` files during ES validation?
A. To enrich data with business context like department or location
B. To store long-term logs for dashboard auditing
C. To provide backup values for CIM field lookup
D. To accelerate the raw data ingestion process

Q9: Which of the following practices helps ensure continuous validation of ES data?
A. Reducing the data retention window for raw logs
B. Running scheduled reports only on summary indexes
C. Setting strict permissions for the admin role
D. Creating test correlation searches to probe data model population

Q10: A correlation search based on `Intrusion_Detection` data model returns no results. What is the most likely cause?
A. The search head is not using REST API mode
B. The user role lacks `schedule_search` capability
C. The TA is missing the `signature` field
D. The data model has no accelerated fields

Learning Path & Study Advice

A strong preparation approach begins with understanding the overall purpose of Splunk Enterprise Security and how it supports security operations beyond basic data search. From there, candidates should build confidence in the platform's main investigative and monitoring workflows, since these reflect how Enterprise Security is used in practice. Once that foundation is established, it becomes easier to study deployment concepts, installation expectations, and the administrative decisions that affect platform readiness.

The next stage of preparation should focus on data quality and context. Candidates benefit from understanding how Enterprise Security depends on validated and normalized data, how add-ons contribute to that process, and why lookups and identity information are essential for interpretation. After that, studying the detection layer becomes more meaningful, especially the distinction between creating correlation searches and tuning them for reliable operational use.

It is also useful to study the platform as an interconnected system rather than as isolated features. Monitoring, investigation, correlation searches, contextual enrichment, visual interfaces, and threat intelligence all work together. A good study approach therefore emphasizes relationships between components, operational purpose, and practical understanding of why each capability matters in a live security environment.

Who This PDF Is For

This document is intended for learners and professionals preparing for responsibilities related to Splunk Enterprise Security administration. It is well suited to security analysts, SOC personnel, SIEM administrators, Splunk administrators moving into security operations work, and technical professionals who support detection and investigation workflows. Readers will benefit most if they already have basic familiarity with Splunk concepts, log data handling, and general cybersecurity principles. It is especially useful for those who want a structured understanding of the certification's knowledge scope without relying on exam-focused shortcuts or memorized fragments.

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-3001 Enterprise Security Administrator Certification Training Course - AAAdemy

Online Flashcards (Quizlet):

https://quizlet.com/user/AAAdemy/folders/splk-3001-splunk-enterprise-security-certified-admin?i=6zfa5t&x=1xqt

Attachment : Answers by Knowledge Point

ES Introduction Practice Question
A1: Answer: A

Explanation: Splunk Enterprise Security (ES) extends the capabilities of Splunk Core by turning it into a full-featured Security Information and Event Management (SIEM) platform. It focuses on detecting, investigating, and responding to cybersecurity threats, which is not the main function of Splunk Core.

A2: Answer: D
Explanation: The Security Posture Dashboard provides a centralized, real-time view of security activity, showing Notable Events, risk scores, and overall security health indicators. It is one of the first views a security analyst sees when logging into Splunk ES.

A3: Answer: C
Explanation: The CIM enables normalization of data from various sources into a consistent format, allowing for standardized searches and correlation rules across different log types. This is essential for effective use of Splunk ES.

A4: Answer: D
Explanation: Notable Events are generated when data matches the conditions defined in correlation searches. These events are highlighted because they represent potential threats that need review or action by a security analyst.

A5: Answer: B
Explanation: The Incident Review Dashboard allows analysts to view, triage, assign, and update the status of Notable Events. It is central to the security operations workflow in Splunk ES.

A6: Answer: A
Explanation: Contextual enrichment refers to the process of enhancing raw events with data such as asset ownership, user identity, department affiliation, or geographic location to provide more useful context during investigations.

A7: Answer: B
Explanation: The Assets and Identities Framework uses lookup tables to map technical identifiers (like IPs and usernames) to real-world context, such as devices or users. This makes event data more meaningful and actionable.

A8: Answer: C
Explanation: Data models in Splunk ES provide a pre-defined schema on top of normalized data, enabling efficient and accelerated searches through mechanisms like tstats, which improves performance and scalability.

A9: Answer: C
Explanation: Correlation searches are rules in Splunk ES that continuously scan data for suspicious patterns. When these patterns are detected, they generate Notable Events for investigation.

A10: Answer: D
Explanation: Splunk ES helps reduce Mean Time to Detect (MTTD) and Mean Time to Respond (MTTR) by offering advanced tools like correlation searches, dashboards, and Notable Events, streamlining the incident lifecycle from detection to resolution.

Monitoring and Investigation Practice Question
A1: Answer: A

Explanation: The Incident Review dashboard is a core feature in Splunk ES used by analysts to manage, filter, and assign Notable Events. It supports team coordination by showing status, urgency, ownership, and event comments.

A2: Answer: B
Explanation: A Notable Event includes contextual fields such as event name, timestamp, source system, risk score, and links to associated logs or users to support investigation.

A3: Answer: D
Explanation: Event Timeline is a visualization tool that helps analysts understand when events occurred and how they relate to each other in time. It enables contextual investigation of incidents.

A4: Answer: C
Explanation: Drilldown and search functionality in Splunk ES allows analysts to click on fields (e.g., IP addresses) and launch deeper SPL-based investigations to explore event context.

A5: Answer: C
Explanation: Analysts should use drilldown and raw log review to evaluate the login behavior, check the risk score and source IP, and determine whether escalation is necessary.

A6: Answer: D
Explanation: Analysts use tags and status indicators to organize and manage Notable Events, helping teams coordinate efforts and maintain investigation histories.

A7: Answer: A
Explanation: Filters in the Incident Review dashboard help security analysts find and prioritize Notable Events by attributes like time, urgency, or assigned analyst.

A8: Answer: B
Explanation: Analysts may close events as false positives when investigation shows the activity is normal, such as expected behavior from a known IP or user.

A9: Answer: B
Explanation: Event Timeline allows security analysts to visually identify how events unfold and correlate over time, providing valuable insight during investigations.

A10: Answer: A
Explanation: Precision investigation in Splunk ES often involves taking a Notable Event, identifying key fields (like IPs or usernames), and launching targeted searches to understand the underlying cause.

Security Intelligence Practice Question
A1: Answer: B
Explanation: Security Intelligence in Splunk ES focuses on reducing alert fatigue and guiding analysts to prioritize alerts that are backed by risk scores and contextual enrichment. It replaces traditional noisy SIEM alerting with smarter, risk-based logic.

A2: Answer: A
Explanation: RBA in Splunk ES works by assigning scores to risky behaviors or events, then triggering a Notable

Event only when the total risk for a user or asset exceeds a certain threshold. This makes alerting more meaningful.

A3: Answer: C
Explanation: The Assets and Identities Framework enriches events by mapping usernames to departments, roles, or importance levels, and IP addresses to known systems. This adds meaningful context to alerts.

A4: Answer: D
Explanation: Risk Rules are correlation searches that assign risk scores to entities instead of creating Notable Events. They contribute to the cumulative risk that may later trigger a single high-priority alert.

A5: Answer: B
Explanation: In RBA, the importance of the asset involved matters. A port scan on a domain controller poses more threat than minor activities on low-value assets like printers or kiosks, so it would carry a higher risk score.

A6: Answer: C
Explanation: Splunk ES aggregates risk scores from multiple risk rules. If the total exceeds a defined threshold, a Notable Event may be created, enabling smarter, cumulative alerting.

A7: Answer: B
Explanation: Contextual enrichment provides key information—such as who the user is, what department they're from, and which asset is involved—so analysts can understand the true significance of an event.

A8: Answer: B
Explanation: The `risk` index stores data from risk rules, including timestamps, risk scores, user and system identifiers, and rule names. This index is critical for RBA functionality.

A9: Answer: A
Explanation: Risk Rules delay alerting by gradually assigning risk scores, which accumulate over time. This reduces alert noise and helps prioritize threats based on context and severity.

A10: Answer: A
Explanation: Security Intelligence leverages correlation, context, and risk scoring to elevate only those incidents that require human attention, reducing noise and alert fatigue in the SOC.

Forensics, Glass Tables, and Navigation Control Practice Question
A1: Answer: C
Explanation: Forensics in Splunk ES is about deeply analyzing security incidents using raw log data, timeline views, and search-based investigation to understand exactly what happened during an attack.

A2: Answer: C
Explanation: Through the Incident Review dashboard, analysts can click into Notable Events and drill down into the raw event logs that triggered them, allowing deeper investigation.

A3: Answer: A
Explanation: Timeline views allow security analysts to understand the chronological flow of events across different systems, which is vital in incident reconstruction and forensics.

A4: Answer: B
Explanation: Glass Tables are designed for executive and managerial audiences who need real-time, high-level visual representations of organizational security health and KPIs.

A5: Answer: A
Explanation: Glass Tables visually display data like risk scores and alerts, often organized by departments or business units, allowing teams to monitor threats in context and act quickly.

A6: Answer: A
Explanation: Navigation Control allows administrators to create role-based user experiences in Splunk ES by adjusting menus, dashboards, and workflows per user or group.

A7: Answer: B
Explanation: By customizing the dashboards and available functions for different roles, Navigation Control makes it easier for new or junior analysts to focus on essential tasks and avoid confusion.

A8: Answer: A
Explanation: Search reconstruction refers to using Splunk searches and filtering to recreate the sequence of events or behaviors during a security incident—essential for forensic analysis.

A9: Answer: D
Explanation: Timeline views provide a visual, chronological layout of related events from multiple sources, helping analysts determine the full scope and order of a potential incident.

A10: Answer: A
Explanation: Glass Tables provide visual context—like color-coded icons and KPIs—so security teams can identify and respond to threats faster based on clearly displayed, up-to-date data.

ES Deployment Practice Question
A1: Answer: A
Explanation: Distributed deployments allocate different components (search heads, indexers, forwarders) across multiple systems, enabling scalability and performance improvements over single-instance setups.

A2: Answer: C
Explanation: Using a dedicated search head for ES isolates the resource-heavy operations of security analytics and correlation searches, preventing interference from other apps and improving performance.

A3: Answer: D
Explanation: SHC provides redundancy and scalability for search heads, allowing multiple analysts to use the system simultaneously without performance degradation.

A4: Answer: C
Explanation: ES versions are tightly bound to specific Splunk Enterprise versions. Reviewing the official compatibility matrix ensures a stable and supported deployment.

A5: Answer: A
Explanation: High-speed disk I/O is crucial for supporting data model acceleration and summary indexing, which are core features in ES that generate heavy read/write activity.

A6: Answer: D
Explanation: The `notable` index is specifically used to store events generated by correlation searches that require review or response from security analysts.

A7: Answer: A
Explanation: In distributed deployments, indexers handle the storage and processing of ingested data, making them essential for search performance and scalability.

A8: Answer: B
Explanation: Keeping security logs in separate indexes makes it easier to perform clean searches, accelerate data models efficiently, and comply with auditing and regulatory requirements.

A9: Answer: B
Explanation: A production-grade ES search head requires at least 16–32 GB of RAM due to high resource usage from dashboards, scheduled searches, and data model acceleration.

A10: Answer: B
Explanation: Splunk ES is resource-intensive, and placing it on a shared search head may lead to degraded performance and reduced effectiveness in running correlation searches.

Installation and Configuration Practice Question
A1: Answer: A
Explanation: In distributed environments, it is common to install Splunk apps via the CLI using `splunk install app` to enable automation and avoid GUI-based limitations.

A2: Answer: D
Explanation: `SA-CIM` is a required supporting app in Splunk ES that enables compatibility with the Common Information Model, ensuring normalized field names across data sources.

A3: Answer: D
Explanation: Post-installation checks include verifying that the navigation menu loads correctly, dashboards display as expected, and data models are populating.

A4: Answer: D
Explanation: DMA precomputes and stores summaries of data to accelerate dashboards and enable real-time correlation searches in ES.

A5: Answer: B
Explanation: The `ess_admin` role has full access to all ES capabilities, including editing correlation searches, risk rules, and managing configurations.

A6: Answer: D
Explanation: Data model acceleration is required for many ES dashboards and for real-time correlation searches to function correctly.

A7: Answer: A
Explanation: These files map IPs and users to human-readable names, departments, and asset priorities, enabling context-rich investigation and alerting.

A8: Answer: B
Explanation: Aggregation policies allow teams to reduce dashboard clutter by combining similar alerts into single incidents for easier triage and investigation.

A9: Answer: A
Explanation: Email alerting allows teams to be notified in real time when high-priority events occur, facilitating faster incident response.

A10: Answer: C
Explanation: Using role inheritance and creating custom roles allows fine-grained access control, improving security and reducing risk of misconfiguration.

Validating ES Data Practice Question
A1: Answer: D
Explanation: The Data Model Audit Dashboard provides insight into how well incoming data aligns with CIM expectations by showing coverage percentage, tagging accuracy, and missing sources.

A2: Answer: A
Explanation: The `| datamodel` command is used to run searches directly against accelerated data models to confirm that the data is populating properly.

A3: Answer: C
Explanation: Data models use tags to categorize events. If the correct tag like `authentication` is missing, events will not be mapped to the model even if they contain the correct fields.

A4: Answer: A
Explanation: CIM requires standardized field names such as `src`, `dest`, etc. Fields like `src_ip` will not be recognized by the data model, so dashboards relying on `src` will not populate correctly.

A5: Answer: C
Explanation: Using the Search Inspector or Field Extractor, analysts can review raw events to confirm whether key CIM fields and tags are properly extracted and visible.

A6: Answer: B
Explanation: If data is being routed to an index that the ES app or user role doesn't have access to, correlation searches will not return results even if they are configured correctly.

A7: Answer: B
Explanation: Event types and tags are required for CIM mapping. Without the correct tags (e.g., `malware`, `authentication`), data will not populate the corresponding data model.

A8: Answer: A
Explanation: These files help enrich log data by associating IPs and users with meaningful context such as department, location, and criticality—essential for risk-based alerting.

A9: Answer: D
Explanation: Creating simple test correlation searches using `| datamodel` is an effective way to confirm whether data is populating the correct CIM-based data models.

A10: Answer: C
Explanation: If the TA (Technology Add-on) or field extractions are missing required fields like `signature`, `src`, or `dest`, the data model won't populate and correlation searches will fail.

Custom Add-ons Practice Question
A1: Answer: A
Explanation: Custom TAs are necessary when vendors don't provide official ones. These add-ons extract and normalize data, tag events appropriately, and map them to CIM so Splunk ES can process them correctly.

A2: Answer: C
Explanation: The Splunk Add-on Builder is the official tool provided by Splunk to help users build custom add-ons with field extractions, event types, and CIM mapping in a guided UI.

A3: Answer: B
Explanation: `src` is the correct CIM-compliant field name used to represent source IP addresses. CIM requires consistent field naming for compatibility with data models.

A4: Answer: D
Explanation: Tags classify event types and are used by Splunk ES to link events to the correct data model. Without proper tags, events may not populate dashboards or trigger correlation searches.

A5: Answer: A
Explanation: Field extractions defined in `props.conf` and `transforms.conf` need to be applied where parsing occurs—either at the heavy forwarder (index time) or on the search head (search time).

A6: Answer: A
Explanation: Modular inputs are commonly used to poll data from external APIs on a scheduled basis. Add-on Builder supports configuring these as part of custom input creation.

A7: Answer: D
Explanation: A TA should focus on data ingestion and normalization, not visualization. Dashboards and saved searches belong in apps, not TAs.

A8: Answer: A
Explanation: The CIM Validation Dashboard in Add-on Builder helps test sample events to ensure they correctly populate CIM data models and meet field/tag requirements.

A9: Answer: B
Explanation: Field aliases help bridge non-standard field names to CIM-compliant ones, improving compatibility when original source fields don't match the CIM format.

A10: Answer: A
Explanation: The Data Model Audit Dashboard allows you to confirm that your TA is successfully feeding data into the expected CIM models. This validates field extraction and tag correctness.

Tuning Correlation Searches Practice Question
A1: Answer: A
Explanation: Tuning correlation searches helps reduce false positives and unnecessary alerts while also minimizing performance impact on the system. This ensures analysts can focus on real threats without overloading the platform.

A2: Answer: C
Explanation: Throttle settings prevent a correlation search from generating duplicate notable events within a defined time window for the same entity (e.g., user, host), keeping dashboards clean and reducing noise.

A3: Answer: C
Explanation: Searching across unnecessarily large time windows on a frequent schedule leads to inefficient searches that consume excessive CPU and memory, especially if the data models are large.

A4: Answer: A
Explanation: Suppression rules are ideal for hiding alerts that result from expected and non-malicious behavior, such as during penetration testing or vulnerability scans, preventing analyst distraction.

A5: Answer: B
Explanation: By adjusting when correlation searches run (e.g., using cron schedules), administrators can control how often detection logic executes, spreading the load and avoiding performance bottlenecks.

A6: Answer: B
Explanation: Assigning risk scores instead of immediate alerts allows events to build contextual weight. Only when thresholds are met is a Notable Event triggered, reducing false positives and emphasizing more serious issues.

A7: Answer: D
Explanation: By adding precise filters (e.g., severity, known assets), you reduce unnecessary matching events and improve the relevance and speed of correlation searches.

A8: Answer: B
Explanation: Adjusting risk thresholds based on business role or asset criticality helps avoid under-alerting for high-value assets and over-alerting for low-impact systems. This makes alerts more actionable.

A9: Answer: A
Explanation: Using cron expressions to space out searches helps avoid peak load situations, especially when multiple high-cost searches would otherwise run simultaneously.

A10: Answer: A
Explanation: Without tuning, correlation searches can flood analysts with alerts, many of which may be false positives. This increases the chance of overlooking actual threats due to alert fatigue.

Creating Correlation Searches Practice Question
A1: Answer: B
Explanation: The Correlation Search Editor in Splunk ES allows configuration of detection settings such as severity, urgency, category, and automated response actions. It connects search logic to workflows.

A2: Answer: A
Explanation: Tokens such as `$user$` are used in correlation search titles or descriptions to dynamically include relevant field values when a notable event is created.

A3: Answer: D
Explanation: Using datamodel-based searches allows you to leverage accelerated data summaries, enabling faster and more efficient searches, especially when working with high-volume authentication data.

A4: Answer: D
Explanation: Running correlation searches in the Search app ensures the search logic behaves as expected, fields are extracted correctly, and results are relevant before introducing it into production workflows.

A5: Answer: C
Explanation: Severity and urgency affect how Notable Events are displayed in dashboards, how analysts prioritize them, and how they influence the overall risk score of users or assets.

A6: Answer: C
Explanation: The Content Management dashboard offers governance for correlation searches and risk rules, including their enablement status, run frequency, and last triggered timestamp.

A7: Answer: C
Explanation: The `lookup` command is used to enhance event data by joining external information, such as watchlists or asset databases, to improve context and detection accuracy.

A8: Answer: A
Explanation: Leveraging `tstats` or datamodels improves performance by querying against accelerated data summaries. This ensures efficiency and scalability of correlation searches.

A9: Answer: A
Explanation: Adaptive response actions can be configured to trigger external systems, such as a ticketing platform, through webhooks or scripted actions integrated with the correlation search.

A10: Answer: B
Explanation: Threat intelligence lookups provide lists of known bad indicators such as IPs or domains. Correlation searches can match observed data against these sources to identify potential threats.

Lookups and Identity Management Practice Question
A1: Answer: C

Explanation: Lookups allow Splunk ES to map values like IP addresses or usernames to enriched contextual data such as departments, asset types, or risk profiles. This enhances alert prioritization and investigation efficiency.

A2: Answer: D
Explanation: The `identities.csv` file maps usernames or email addresses to contextual identity data including departments, names, and access tiers. It's a critical part of Splunk ES identity enrichment.

A3: Answer: A
Explanation: KV Store lookups are designed for dynamic data that can be updated via the Splunk Web interface or scripts. They support real-time adds, edits, and deletions, making them ideal for managing live asset data.

A4: Answer: D
Explanation: External lookups are script-based and can be used to fetch data from external services like VirusTotal or a threat intelligence API. They provide real-time enrichment for advanced context in detection logic.

A5: Answer: B
Explanation: The `assets.csv` and `identities.csv` files are managed under the `SA-IdentityManagement` app, specifically in its `lookups` directory. This supports identity context used across Splunk ES.

A6: Answer: B
Explanation: The `priority` field in asset or identity lookups influences how much risk weight is assigned to an event involving that entity. Values such as Critical, High, or Low help tailor alert significance.

A7: Answer: A
Explanation: The Asset and Identity Center is a dashboard in Splunk ES that allows users to view and validate loaded identity/asset data, ensuring mappings are correct and properly used in searches.

A8: Answer: C
Explanation: By using lookups in correlation searches, you can enrich detections with role, department, priority, or asset criticality, enabling more targeted and useful alerting.

A9: Answer: A
Explanation: A typical CSV lookup is formatted using comma-separated values where headers define field names. Each row maps values used in lookups (e.g., IP to hostname).

A10: Answer: B
Explanation: KV Store lookups are editable at runtime using the UI or REST API, allowing security teams to quickly update identity or threat data without requiring a Splunk restart or re-upload.

Threat Intelligence Framework Practice Question
A1: Answer: C
Explanation: The Threat Intelligence Framework ingests and manages threat intelligence feeds, then matches known indicators (such as malicious IPs or file hashes) against log data to help identify threats in real-time.

A2: Answer: A
Explanation: TIF natively supports the TAXII protocol, which is designed for the automated exchange of

STIX-formatted threat intelligence. This makes it a preferred choice for integration with government or industry threat-sharing groups.

A3: Answer: C
Explanation: Splunk ES organizes threat intelligence into dedicated lookup tables. File hash indicators are stored in `threat_intel_by_file_hash`, which can be queried to detect file-based threats.

A4: Answer: C
Explanation: KV Store collections are used in TIF to store threat indicators (IOCs), manage expiration policies (aging), and add enrichment tags such as source, threat type, and confidence.

A5: Answer: D
Explanation: The Threat Intelligence Manager dashboard helps analysts track the health of their threat feeds, verify which IOCs have been ingested, and monitor IOC match rates and coverage.

A6: Answer: C
Explanation: The `confidence` field is attached to IOCs and helps analysts determine how reliable the threat data is. Values like "low", "medium", and "high" assist in triaging alerts.

A7: Answer: B
Explanation: Threat Match Search compares live or recent log data against loaded threat intelligence to find matches with IOCs (IPs, domains, hashes, emails) and initiate triage or response.

A8: Answer: A
Explanation: TIF adds valuable threat intelligence metadata to notable events, such as threat type and confidence, which helps analysts prioritize incidents and understand the context faster.

A9: Answer: A
Explanation: TIF allows manual uploading of IOC data using structured formats such as CSV or JSON. These files can include fields like `ip`, `threat_type`, and `confidence`.

A10: Answer: A
Explanation: Many IOC records in KV Store use a TTL (time to live) setting to automatically expire after a configured time period. This prevents stale indicators from polluting correlation searches.