Up-to-date Practice Test with Latest Questions and Answers covering latest syllabus and topics of the exam. Makes you ready to face actual exam.

SPLK-1005 Practice Questions
SPLK-1005 Practice Test
SPLK-1005 Practice Exam
SPLK-1005 Exam Questions
SPLK-1005 Study Guide

killexams.com

Splunk

SPLK-1005

Splunk Cloud Certified Admin

ORDER FULL VERSION

https://killexams.com/pass4sure/exam-detail/SPLK-1005

Question: 999

For a monitor input collecting data from /var/log/apache2/ with mixed access.log and error.log files requiring different sourcetypes, which props.conf configuration correctly assigns sourcetypes by source?

A. [source::/var/log/apache2/access.log] sourcetype = apache_access [source::/var/log/apache2/error.log] sourcetype = apache_error
B. [source::apache2] sourcetype = auto_sourcetype regex = (access|error)
C. [sourcetype::access_log] TRANSFORMS = apache_access [sourcetype::error_log] TRANSFORMS = apache_error
D. [apache_access] TIME_FORMAT = %d/%b/%Y:%H:%M:%S [apache_error] TIME_FORMAT = %Y-%m-%d

Answer: A

Explanation: The correct props.conf configuration uses source stanza format [source::] to assign sourcetypes based on the full file path. Each log file gets its own stanza with explicit sourcetype assignment: access.log maps to apache_access and error.log maps to apache_error. This approach provides precise control over sourcetype assignment during the input phase, ensuring proper parsing rules are applied to each log type.

Question: 1000

When configuring props.conf for anonymization, the DEST_KEY parameter in transforms.conf controls where the transformed data is written. Which DEST_KEY value correctly writes the masked result back to replace the original raw event data?

A. DEST_KEY = _meta
B. DEST_KEY = _sourcetype
C. DEST_KEY = _raw
D. DEST_KEY = _time

Answer: C

Explanation: DEST_KEY = _raw writes the transformed data back to the raw event field, effectively replacing the original raw data with the masked version. This is required for data anonymization and masking operations where the original sensitive data must be replaced in the indexed events.

Question: 1001

An organization requires account lockout for suspected brute-force attempts when using LDAP authentication against Active Directory in Splunk Cloud. Which approach will produce a compliant, Splunk Cloud-compatible solution?

A. Configure local Splunk user lockouts and synchronize lockout state back to AD via LDAP write-back
B. Configure account lockout policies on AD and enable LDAP bind user limited retries, while enforcing TLS between Splunk and AD
C. Add authLockout = true in authentication.conf on Splunk Cloud to enable lockouts locally
D. Use SAML instead and rely on IdP lockout settings while keeping anonymous LDAP failsafe enabled

Answer: B

Explanation: Splunk Cloud delegates account lockout behavior to the external directory for LDAP/AD authentication; enforce TLS for secure binds and configure AD account lockout policies server-side. There is no supported configuration to centrally enable lockout inside

Splunk Cloud for LDAP accounts — AD-side policies and secure LDAP binds are the appropriate control.

Question: 1002

An organization wants to send system logs from their public cloud infrastructure directly into Splunk Cloud via the HTTP Event Collector (HEC). They have a Splunk Cloud Managed deployment. Which step must the Splunk Cloud Admin perform to enable this ingestion path?

A. Edit the inputs.conf file on the Splunk Cloud Search Head Cluster to listen on a custom UDP port.
B. Generate a client-side SSL certificate on an on-premises certificate authority and upload it to the Splunk Cloud platform via an enhancement ticket.
C. Create a HEC token via the Splunk Cloud Web UI or the Admin Config Service (ACS) API, and ensure that the HEC token is enabled and targeted at the appropriate HEC endpoint provided in their Splunk Cloud details.
D. Open port 8088 on the Splunk Cloud indexers by modifying the network security groups inside the cloud provider's console.

Answer: C

Explanation: In Splunk Cloud, the HTTP Event Collector (HEC) is managed through the native Splunk interfaces or the Admin Config Service (ACS) API. The Admin generates a HEC token, assigns it to specific indexes, and ensures it is enabled. Splunk Cloud provides specific URLs for HEC traffic (often using port 443 or 8088 depending on the configuration). Network access control and port provisioning are handled automatically by the platform; the customer does not modify cloud provider network security groups.

Question: 1003

A Windows system administrator needs to configure a Universal Forwarder to collect

security events from the local Windows Event Log system. Which input stanza type must be used within the inputs.conf file to natively ingest these specialized operating system events?

A. [wmi_monitor://security]
B. [WinEventLog://Security]
C. [wnt_event://security_log]
D. [windows_log://security]

Answer: B

Explanation: Splunk utilizes specialized input types to interact natively with Windows operating system logging APIs. The standard syntax to ingest local Windows Event Logs requires the [WinEventLog://] stanza format, where matches the event log channel like Security or System.

Question: 1004

During the ingestion of application traces, a Splunk administrator notices that multi-line stack traces are being fragmented into multiple single-line events. The logs contain a distinct timestamp at the start of each valid log entry formatted as YYYY-MM-DD HH:MM:SS. Which configuration settings should be applied in props.conf on the parsing tier to ensure that these multi-line logs are assembled into unified events?

A. [app_traces]
BREAK_ONLY_BEFORE = ^\d{4}-\d{2}-\d{2} \d{2}:\d{2}:\d{2}
MAX_EVENTS = 500
SHOULD_LINEMERGE = false
B. [app_traces]
BREAK_ONLY_BEFORE_DATE = true
MAX_EVENTS = 500
SHOULD_LINEMERGE = true
C. [app_traces]
MUST_BREAK_AFTER = ^\d{4}-\d{2}-\d{2} \d{2}:\d{2}:\d{2}
MAX_EVENTS = 500

SHOULD_LINEMERGE = true

D. [app_traces]
TIME_PREFIX = ^
MAX_EVENTS = 500
SHOULD_LINEMERGE = false

Answer: A

Explanation: Splunk best practices for high-performance multi-line event processing dictate setting SHOULD_LINEMERGE = false and utilizing BREAK_ONLY_BEFORE along with an explicit regular expression matching the event start pattern. This approach instructs the line-breaking engine to break data streams into events more efficiently than using the older, resource-intensive SHOULD_LINEMERGE = true line merging architecture, while the MAX_EVENTS setting ensures that excessively long stack traces do not exceed the processing window limits.

Question: 1005

An enterprise organization wants to enforce multi-factor authentication (MFA) for its local Splunk Cloud accounts that do not use the corporate single sign-on (SSO) system. The security policy dictates the use of Time-based One-Time Password (TOTP) applications like Google Authenticator. Which administrative action must be taken first to enable and enforce this configuration across the Splunk environment?

A. Modify the local configuration file authentication.conf via an SSH session to configure the internal multi-factor parameters.
B. Navigate to Settings > Authentication Methods, select Multifactor Authentication, choose Time-based One-Time Password, and configure it.
C. Install a dedicated premium security application from Splunkbase that overwrites the default user login validation mechanisms.
D. Navigate to Settings > Authentication Methods > Duo Security, and insert the registration integration keys provided by the vendor.

Answer: B

Explanation: Splunk Cloud natively supports native Time-based One-Time Password (TOTP) multi-factor authentication for local accounts. This can be enabled by an administrator through the Splunk Web interface by navigating to Settings > Authentication Methods, selecting the Multifactor Authentication radio button, picking TOTP from the provider options, and completing the configuration steps. Direct file modification via SSH is not available in Splunk Cloud.

Question: 1006

An administrator configures an inputs.conf monitor stanza to watch a directory for incoming CSV reports. The files are generated daily and named using the convention report_YYYYMMDD.csv. The administrator wants to ensure that if an older file is updated for auditing purposes, Splunk will re-read the entire file from the beginning rather than just trying to index the appended lines. Which attribute configuration achieves this behavior?

A. initCrcLength = 1024
B. crcSalt =
C. always_readme = true
D. update_type = reset

Answer: B

Explanation: By default, Splunk tracks monitored files using a hash of their initial 256 bytes stored in the fishbucket database. If a file name changes or data is modified but the first 256 bytes remain identical, Splunk assumes it is the same file and looks for an offset pointer. Setting crcSalt = forces Splunk to include the unique file path string as part of the CRC hash generation, ensuring that different file names or re-created files with identical headers are treated as entirely unique entities and re-indexed from the start.

Question: 1007

An administrator needs to exclude a specific set of staging servers from a server class that currently targets an entire subnet. The server class definition uses a whitelist for the subnet. Which configuration syntax in serverclass.conf correctly implements the exclusion of hosts named 'stage-app01' through 'stage-app09'?

A. blacklist.0 = stage-app0?
B. blacklist.0 = stage-app0[1-9]*
C. exclude.0 = stage-app0[1-9]
D. blacklist.1 = stage-app0[1-9]

Answer: D

Explanation: In serverclass.conf, filters are evaluated in a specific sequence where blacklists take precedence over whitelists. The correct attribute name is blacklist followed by an index number (e.g., blacklist.1). The value supports standard Splunk wildcarding and character classes, meaning stage-app0[1-9] will precisely match and exclude stage-app01 through stage-app09 from the server class.

Question: 1008

A Splunk Cloud customer wants to prevent specific IdP username values from logging in (for example, service accounts that were disabled in the IdP but still present). Which Splunk configuration enforces this exclusion on the Splunk side for SAML-authenticated users?

A. Set ignoreUnknownUsers = true to drop unknown SAML accounts
B. Configure excludedUsers in authentication.conf to list the usernames to prevent login
C. Configure excludedSAMLGroups to list the groups that contain those accounts
D. Modify roles.conf to set allowLogin=false for those usernames

Answer: B

Explanation: The excludedUsers setting in authentication.conf can be used to specify a

comma-separated list of user names from the SAML response that Splunk should exclude from logging in, preventing those IdP usernames from accessing Splunk even when they are present in SAML assertions.

Question: 1009

A Splunk Cloud Admin wants to optimize search performance across a large cluster. They observe that user queries frequently perform wildcard searches at the beginning of fields (e.g., *value). The Admin decides to enable segment term indexing to speed up these leading wildcard queries. Which statement correctly describes how this task is executed in Splunk Cloud?

A. The Admin must log into the Splunk Cloud indexers via SSH using their admin credentials and append the configuration to /opt/splunk/etc/system/local/indexes.conf.
B. The Admin can navigate to Settings > Indexes in the Splunk Cloud Web UI, select the target index, and toggle the "Enable Leading Wildcard Optimization" checkbox.
C. The Admin cannot directly modify indexes.conf on the indexers; they must use the Admin Config Service (ACS) API or submit a Splunk Support ticket to configure segment term indexing (INDEXED_VALUE = true or specialized segmenters) for specific indexes.
D. The Admin must deploy an app containing an updated segmenters.conf file directly to the search head cluster using the forwarder management console.

Answer: C

Explanation: In Splunk Cloud, direct command-line access to the underlying infrastructure (such as SSH access to indexers) is prohibited. Configurations that alter index structures or index-time properties like segment term indexing reside in indexes.conf or segmenters.conf. Because these changes impact indexer resources and storage, the Admin must use supported cloud administration interfaces like the Admin Config Service (ACS) API or work through Splunk Support to implement these deep index-tier modifications safely.

Question: 1010

An enterprise deployment uses an automated script to export database tables into flat text files within a monitored directory /data/exports/. The files are named dynamically based on the table name, such as users_export.txt or inventory_export.txt. The administrator wants to configure a single monitor stanza that captures all text files in this directory but ensures that any file containing the string "test" or "backup" anywhere in the path is completely skipped. Which configuration satisfies this requirement?

A. [monitor:///data/exports/.*\.txt]
blacklist = .*test.*|.*backup.*
B. [monitor:///data/exports/*.txt]
ignoreFiles = *test*,*backup*
C. [monitor:///data/exports/.../*.txt]
exclude = test,backup
D. [monitor:///data/exports]

filter = negate(test,backup)

Answer: B

Explanation: The [monitor:///data/exports/*.txt] stanza targets all files with a .txt extension within the specified directory path. The blacklist attribute takes a regular expression string; using .*test.*|.*backup.* evaluates the file path and accurately flags any file containing "test" or "backup" as excluded, matching the administrative requirement.

Question: 1011

Which deployment-server scaling practice is most appropriate when managing several thousand forwarders?

A. Segment forwarders into multiple deployment servers based on geographic region or business unit.
B. Run a single, large deployment server instance and rely on built-in queuing to handle all clients.

C. Replace the deployment server with a cloud-native configuration service that pushes configs via API.
D. Use multiple deployment servers in a load-balanced cluster fronted by a DNS round-robin.

Answer: A

Explanation: When managing several thousand forwarders, segmenting them across multiple deployment servers by region or business unit reduces the load on any single server and improves manageability. Each deployment server can independently manage its own pool of clients, while policy and app definitions can still be kept consistent through shared configuration artifacts.

Question: 1012

A team wants SSO through an identity provider and wants Splunk to trust assertions from that provider rather than a password stored in Splunk. Which authentication method is the best fit?

A. LDAP authentication, because it replaces password validation with directory lookup.
B. Scripted authentication, because it is the standard method for all modern SSO deployments.
C. Native Splunk authentication, because it stores credentials locally and supports SSO claims.
D. SAML authentication, because it uses assertions from an identity provider for login.

Answer: D

Explanation: SAML is the correct fit when the identity provider performs authentication and Splunk consumes the resulting assertion. This supports single sign-on patterns where Splunk trusts the external identity provider. LDAP is directory-based authentication, but SAML is the typical modern SSO model.

Question: 1013

A security team requires that data in the audit_logs index be securely and permanently destroyed as soon as it leaves the active search tiers after 180 days. They want to ensure no remnants are left on the local storage systems. What is the default behavior of Splunk Cloud when data reaches its frozen criteria, and how satisfies this requirement?

A. Splunk Cloud automatically archives all frozen buckets to an internal glacier tier unless a custom script is uploaded via the App Management page.
B. Splunk Cloud moves frozen data to a recycled directory that must be cleared manually by raising a support ticket.
C. Splunk Cloud executes a secure DoD-compliant multi-pass overwrite on the underlying disk blocks whenever a bucket rolls to frozen.
D. Splunk Cloud automatically deletes the buckets from the file system permanently when they transition to the frozen state, unless an archiving destination is explicitly configured.

Answer: D

Explanation: By default, when an index bucket satisfies the conditions to transition to the frozen state (either via exceeding frozenTimePeriodInSecs or maxTotalDataSizeMB), Splunk Cloud permanently deletes the bucket from the storage system. Data is only preserved if the administrator has explicitly configured self-storage archiving to export the frozen buckets to a customer-managed cloud storage bucket.

Question: 1014

An administrator wants to view the current status of all deployment clients, including their last check-in times and applied server classes, using the Splunk Command Line Interface (CLI) on the deployment server. Which CLI command provides this specific administrative information?

A. splunk show deployment-status

B. splunk status server-class
C. splunk display deployment-clients
D. splunk list deploy-clients

Answer: D

Explanation: The splunk list deploy-clients command is a built-in Splunk CLI utility executed on the deployment server. It queries the internal deployment server state and lists details for all registered deployment clients, including their hostnames, IP addresses, operating systems, last phone-home timestamps, and assigned server classes.

Question: 1015

A Splunk Admin needs to parse data where a single physical line contains multiple distinct events jammed together, separated by the character sequence ###. The administrator wants to split this single line into multiple separate index events. Which attribute in props.conf should be configured to split the stream based on this delimiter?

A. [multi_event_line]
BREAK_ONLY_BEFORE = ###
B. [multi_event_line]
LINE_BREAKER = (###)
C. [multi_event_line]
LINE_BREAKER = (?:###)+
D. [multi_event_line]
MUST_BREAK_AFTER = ###

Answer: B

Explanation: The LINE_BREAKER attribute in props.conf accepts a regular expression that must contain a capturing group. When the ingestion engine evaluates the data stream, it finds matches for the pattern, discards the text captured by the capturing group, and splits the stream into a new event at that boundary. This occurs prior to line merging or standard event

processing, making it the most efficient way to break single physical lines containing multiple logical events into individual entries.

Question: 1016

A security administrator wants to restrict access to the Splunk Cloud Web UI so that it can only be accessed from the corporate office network and the company's authorized VPN IP ranges. How should this configuration be implemented in a standard Splunk Cloud environment?

A. Use the Admin Config Service (ACS) API or the Splunk Cloud Web UI to add the designated IP CIDR blocks to the Search Head access list.
B. Open a support ticket to modify the underlying AWS Security Groups associated with the load balancers.
C. Deploy an on-premises firewall proxy that rewrites the HTTP headers of the incoming requests to match internal Splunk security tokens.
D. Modify the web.conf file on the Search Head Cluster Deployer and add the IP addresses to the mgmtHostUris parameter.

Answer: A

Explanation: Splunk Cloud allows administrators to manage their own network security perimeters via IP allowlists. This can be configured directly through the Admin Config Service (ACS) API or via the Splunk Web UI under the appropriate administrative configurations for traffic types like Search Head access, HEC, or forwarder ingestion. Direct modification of web.conf or underlying AWS Security Groups by the customer is not allowed.

Question: 1017

A heavy forwarder acts as an intermediate aggregation layer, accepting raw logs over a TCP port from a router farm. The incoming logs vary dynamically in character encoding, with

some legacy routers sending data in Latin-1 (ISO-8859-1) and newer routers sending data in UTF-8. Which configuration parameter in inputs.conf can be set at the stanza level to force Splunk to validate and fallback gracefully when non-UTF-8 characters are received?

A. [tcp://5142] \n connection_host = ip \n encoding = UTF-8
B. [tcp://5142] \n disabled = 0 \n fallback_charset = iso-8859-1
C. [tcp://5142] \n disabled = 0 \n route = has_header
D. [tcp://5142] \n charset = latin1 \n disabled = 0

Answer: A

Explanation: Splunk provides the encoding setting within inputs.conf for network inputs. Setting encoding = UTF-8 or specific target encodings tells Splunk how to interpret incoming byte streams. If the incoming bytes do not conform strictly to UTF-8, Splunk applies a lossy conversion or fallback decoding mechanism to convert invalid sequences into valid characters, ensuring that the indexing process does not crash or corrupt the indexing pipeline.

Question: 1018

A Splunk administrator needs to provision an app that contains sensitive database credentials to a subset of heavy forwarders. Because the deployment server transmits apps unencrypted over the management port by default, what is the Splunk-recommended practice to secure these credentials during transit to deployment clients?

A. Use a script to obfuscate the credentials using Splunk's internal secret key before deploying the app
B. Enable SSL/TLS verification by configuring the sslVerifyServerCert attribute in deploymentclient.conf
C. Encrypt the app payload using a local zip password before placing it in the deployment-apps directory
D. Configure the targetRepository stanza to route traffic through an alternative SSH tunnel on port 22

Answer: B

Explanation: To secure data in transit between the deployment server and deployment clients, administrators should enable full SSL validation. Configuring sslVerifyServerCert and specifying valid certificates in deploymentclient.conf and server.conf ensures that the communication channel over the management port is encrypted and verified against trusted certificate authorities.

KILLEXAMS.COM

Killexams.com is a leading online platform specializing in high-quality certification exam preparation. Offering a robust suite of tools, including Exam Questions, practice tests, and advanced test engines, Killexams.com empowers candidates to excel in their certification exams. Discover the key features that make Killexams.com the go-to choice for exam success.

Practice Exam Questions Based on Current Exam Objectives

Killexams.com provides practice exam questions aligned with the latest official exam objectives and latest syllabus. Our content is reviewed and updated regularly to reflect recent changes announced by certification vendors. By studying these practice questions, candidates will cover the structure, difficulty level, and topics of the actual exam, helping them prepare more effectively and efficiently.

Comprehensive Practice Exams (PDF Format)

Killexams.com offers multiple-choice questions (MCQs) in easy-to-read PDF format, covering all major domains of the exam. Each PDF contains a structured collection of practice questions and verified answers designed to support focused study. These MCQs help candidates reinforce key concepts, identify knowledge gaps, and improve exam readiness through consistent practice.

Realistic Practice Tests (Online Test Engine & Desktop Test Engine)

To support hands-on preparation, Killexams.com provides practice tests through both an Online Test Engine and a Desktop Test Engine. These tools are designed to simulate a real exam environment, allowing candidates to practice under exam-like conditions, with latest syllbus and topics of the exam. Performance tracking, test history, and result analysis help users evaluate their progress and focus on areas that need improvement.

Risk-Free Purchase Policy

Killexams.com follows a transparent and customer-friendly purchase policy. If users are not satisfied with the study materials, they may request assistance or a refund in accordance with our published terms and conditions. This policy reflects our commitment to customer satisfaction, fairness, and confidence in our preparation resources.

Regularly Updated Content

Our practice question bank is reviewed and updated on an ongoing basis to stay aligned with the latest exam outlines and vendor updates. This ensures candidates are studying up-to-date, relevant material, and preparing with content that reflects current exam expectations, helping them stay confident and well-prepared.