**Exam** : **SPK-503**

**Title** : Splunk Certified Cybersecurity Defense Architect

**Version** : DEMO

1.AJ has been assigned to design controls for a new, low-latency and highly resilient application. The business requires no downtime if a device fails or while maintenance is performed.
Which of the following deployment options will satisfy these requirements?
A. Active/passive cluster
B. Active/active cluster
C. Diffused cluster
D. Distributed cluster
**Answer:** B
**Explanation:**
An active/active cluster is designed to maintain multiple operational service instances that process workloads simultaneously, providing redundancy and supporting continuous service availability.
In an active/active architecture, traffic can be distributed across healthy nodes using load-balancing mechanisms. If one node fails or requires maintenance, the remaining nodes can continue processing requests without waiting for a passive replacement to become operational.
This makes active/active clustering particularly appropriate for applications requiring low latency, resilience, and minimal interruption.
For example, a security gateway deployment supporting a high-volume application may distribute requests across several active instances. When a node becomes unavailable, traffic can be redirected to the remaining healthy instances.
However, achieving genuinely uninterrupted operation also requires appropriate session handling, connection draining, health checks, capacity planning, and redundant supporting infrastructure.
An active/passive cluster provides redundancy but can introduce a failover interval while the standby node assumes responsibility.
A distributed architecture may improve scalability, but distribution alone does not guarantee continuous service during failures.
A diffused cluster is not a conventional high-availability deployment model in this context. Consequently, active/active clustering most directly satisfies the requirement for resilient, low-latency processing during device failures and planned maintenance.
Official reference /study topic: Splunk certification blueprint, Security Capability Selection: high availability and resilience.

2.The increasing rate of cyberattacks has caused many countries to enact laws and regulations concerning the collection, handling, and protection of their citizens' private information, regardless of where it is stored.
Which term best describes these laws?
A. Data residency
B. Data sovereignty
C. Data shielding
D. Data providence
**Answer:** B
**Explanation:**
Data sovereignty concerns the application of legal and regulatory authority to data, including requirements governing its collection, processing, access, protection, storage, and transfer.
The concept is particularly important for multinational organizations because information may be subject

to legal obligations arising from the location of individuals, processing activities, service providers, or applicable jurisdictions.
For example, GDPR can apply to organizations outside the European Union when their activities meet specified territorial-scope conditions.
An organization cannot assume that relocating personal information to infrastructure in another country automatically removes applicable privacy obligations.
Data sovereignty considerations therefore influence cloud architecture, contractual arrangements, cross-border transfers, government-access requirements, and the selection of service providers.
Data residency, option A, refers more specifically to the geographic location where information is stored or processed.
Residency requirements may be one mechanism used to satisfy applicable legal obligations, but physical location and legal jurisdiction are not identical concepts. Data shielding is not the standard term for the complete set of laws described. Data providence is not the recognized legal concept in this context.
It is also important to distinguish data sovereignty from an unrestricted claim that every country's laws automatically apply to all information concerning its citizens everywhere. Legal applicability depends on the relevant statute and its jurisdictional provisions. Official reference /study topic: Splunk Cybersecurity Defense Architect: Governance Risk and Compliance, Data Sovereignty and Residency.

3.Which type of data does OpenTelemetry supply that the security team can use during an investigation?
A. SBOM
B. Threat path
C. Normalization
D. Traces
**Answer:** D
**Explanation:**
OpenTelemetry is an open-source observability framework that enables applications and infrastructure to generate, collect, process, and export telemetry, including traces, metrics, and logs.
Distributed traces provide visibility into how individual requests move through interconnected application components and services.
A trace generally consists of related spans representing operations performed during a request's execution.
For example, a customer-facing transaction may generate spans associated with an API gateway, authentication service, application server, database, and downstream processing components.
During a security investigation, these traces can help analysts reconstruct execution paths, identify unusual service interactions, and correlate suspicious application behavior with other telemetry.
They can also reveal dependencies between services and help distinguish application failures from potentially unauthorized activities.
The OpenTelemetry specification defines traces through relationships between spans, enabling distributed request activity to be correlated across processing boundaries.
A Software Bill of Materials identifies software components and dependencies rather than runtime request execution.
Threat paths are typically derived from security graph analysis, exposure management, or adversary-path modeling.

Normalization refers to standardizing data representation and is not a distinct OpenTelemetry telemetry signal.
Therefore, traces are the appropriate selection.
OpenTelemetry traces should be correlated with identity, network, endpoint, and application security events to provide stronger investigative evidence. Official reference /study topic: OpenTelemetry, Signals and Tracing Signal
Specification.

4.Which MLTK command can be used with tstats in an ES detection to apply a machine learning model to the search results?
A. Summary
B. Fit
C. Cluster
D. Sample
**Answer:** B
**Explanation:**
The Splunk Machine Learning Toolkit (MLTK) provides specialized Search Processing Language commands for applying statistical and machine-learning algorithms to security telemetry. Among the listed options, fit is the appropriate command because it trains a machine-learning model and applies that model to the current search results.
The tstats command efficiently retrieves aggregated information from indexed fields and accelerated data models. It can produce datasets containing activity frequencies, authentication counts, network connections, and other security measurements. These results can be passed into supported machine-learning algorithms for behavioral analysis and anomaly identification.
An example workflow is:
| tstats count where index=security by host
| fit DensityFunction count
The fit operation evaluates the resulting data and generates model-derived output for analysis.
However, an important distinction exists in Splunk's official documentation: ‘apply` is the dedicated command for applying a previously trained and saved model to new search results. That command is missing from the supplied choices. Therefore, B is the best available answer, but the question is technically imprecise if it specifically refers to an existing model.
summary describes a learned model, sample selects subsets of data, and cluster is not the MLTK operation being requested.
Official reference /study topic: Splunk Machine Learning Toolkit, Search Commands for Machine Learning: fit and apply; certification blueprint Objective 4.4.

5.How should the control network in a water treatment plant be segregated from other networks?
A. With Access Control Lists
B. Physically, using a diode if needed
C. Physically, using a VPN if needed
D. Using RFC1918 addresses
**Answer:** B
**Explanation:**

Water treatment plants operate critical Operational Technology (OT) environments where cybersecurity failures can affect physical processes, public safety, and service availability. Physical segregation provides a strong architectural boundary between sensitive industrial control systems and less-trusted enterprise or external networks.

Where operational requirements permit, separate network infrastructure can reduce opportunities for attackers to move from compromised corporate systems into supervisory control, process automation, and safety-related environments.

A unidirectional gateway, commonly called a data diode, provides an additional protective mechanism by allowing information to flow in only one permitted direction.

For example, operational telemetry may be exported from an isolated control environment to a monitoring network without permitting inbound command traffic through the same connection.

NIST SP 800-82 Rev. 3 discusses physical and logical segmentation, network zones, industrial demilitarized zones, firewalls, and unidirectional gateways as components of OT defense-in-depth architecture.

Physical segregation is not universally mandatory for every OT connection. The selected architecture must account for required data flows, operational safety, maintenance access, and risk assessment.

Access Control Lists provide logical restrictions but do not necessarily establish equivalent isolation.

VPNs protect communications over networks but do not independently create physical segregation.

RFC1918 addresses merely designate private IPv4 address ranges. Therefore, option B provides the strongest segregation method among the listed choices. Official reference /study topic: NIST SP 800-82 Rev. 3, Guide to Operational Technology Security, segmentation and isolation.