Exam : **SPK-202**

Title : Splunk Enterprise Certified Architect

Version : DEMO

1.A Splunk Enterprise project has an agreed charter describing its purpose, scope and stakeholders. The team has also validated its initial use cases and data sources.
Which additional document most directly connects those requirements to implementation?
A. A statement of the high-level goals and stakeholder responsibilities.
B. An inventory of current hosts, installed software and support contacts.
C. A plan relating the validated requirements to architecture, prerequisites and implementation steps.
D. An operating guide describing routine monitoring, escalation and service ownership.
Answer: C
Explanation:
An implementation plan turns the validated requirements into a design the team can build. It ties each use case and data source to the target architecture, its prerequisites and the steps needed to deploy it.

2.An organization is building a new Splunk Enterprise 9.3 deployment with one search head, multiple separate indexers and forwarders. It is not migrating an existing instance. The implementation roadmap and the handoff after components begin communicating are being reviewed.
Which proposal correctly addresses BOTH decisions?
A. Use the single-search-head/multiple-indexer framework; then configure users, roles, security and useful knowledge objects.
B. Use the search-head-cluster framework; then configure users, roles, security and useful knowledge objects.
C. Use the single-search-head/multiple-indexer framework; treat user, role and security setup as already completed by establishing component communication.
D. Use the search-head-cluster framework; treat user, role and security setup as already completed by establishing component communication.
Answer: A
Explanation:
The single-search-head framework matches this deployment: one search head queries several indexers. Once the components communicate, the team can configure users, roles and security, then create the knowledge objects needed for the use cases.

3.A Splunk Enterprise capacity worksheet already records 110 GB/day of incoming data, 280 registered users, the deployment hardware, search types, saved-search schedules and the apps to be installed. Two entries are still missing: the amount of searchable data that must accumulate, and the peak simultaneous search load.
Which TWO questions most directly fill those gaps? (Choose two)
A. What is the average daily incoming data volume?
B. How long must each relevant data set remain searchable?
C. How many user accounts are registered?
D. Which apps will be installed on the search tier?
E. How many users and searches will be active concurrently at peak?
Answer: B, E
Explanation:
Retention determines how many days of each data set accumulate at the expected daily volume. Peak concurrent users and searches determine how much search work the system must support at the same

time.

4.A team says its proposed Splunk Enterprise hosts meet the installation minimums and asks the architect to approve capacity. The team has not documented expected data growth, search-response objectives, access to source data or staffing constraints.
What is the most appropriate next planning action?
A. Approve the design from the installation minimums and review growth after production use begins.
B. Collect the business and workload constraints, then assess reference hardware against the expected load and growth.
C. Replace the workload assessment with a check that every host uses the same hardware specification.
D. Approve the design from current daily ingestion, then treat user activity as a separate operating concern.
Answer: B
Explanation:
Data access, staffing and implementation constraints determine whether the design is feasible. Expected volume, arrival rate, performance targets and growth give the architect a basis for sizing the hardware.

5.An architect is applying the Splunk Enterprise 10.4 performance questionnaire to its stated single-instance reference-machine starting point. The proposed workload is 250 GB/day with six concurrent users.
Considering questions 3 and 4 specifically, which conclusion follows?
A. Keep one dedicated instance because ingestion is below 300 GB/day, regardless of the concurrent users.
B. Keep one dedicated instance because distribution requires both ingestion and concurrent-user thresholds to be exceeded.
C. Distribute across multiple machines because the concurrent-user branch is exceeded, while retaining the questionnaire's reference-machine context.
D. Reject the workload because more than four concurrent users exceeds Splunk Enterprise's supported product limit.
Answer: C
Explanation:
For the questionnaire's reference machine, a requirement above four concurrent users calls for distribution across multiple machines. Six concurrent users meet that condition.