

IT 认证电子书



质 量 更 高 服 务 更 好

半年免费升级服务

<http://www.itrenzheng.com>

Exam : **SPLK-2002**

Title : Splunk Enterprise Certified
Architect Exam

Version : DEMO

1.Which of the following will cause the greatest reduction in disk size requirements for a cluster of N indexers running Splunk Enterprise Security?

- A. Setting the cluster search factor to N-1.
- B. Increasing the number of buckets per index.
- C. Decreasing the data model acceleration range.
- D. Setting the cluster replication factor to N-1.

Answer: C

Explanation:

Decreasing the data model acceleration range will reduce the disk size requirements for a cluster of indexers running Splunk Enterprise Security. Data model acceleration creates tsidx files that consume disk space on the indexers. Reducing the acceleration range will limit the amount of data that is accelerated and thus save disk space. Setting the cluster search factor or replication factor to N-1 will not reduce the disk size requirements, but rather increase the risk of data loss. Increasing the number of buckets per index will also increase the disk size requirements, as each bucket has a minimum size. For more information, see Data model acceleration and Bucket size in the Splunk documentation.

2.Stakeholders have identified high availability for searchable data as their top priority.

Which of the following best addresses this requirement?

- A. Increasing the search factor in the cluster.
- B. Increasing the replication factor in the cluster.
- C. Increasing the number of search heads in the cluster.
- D. Increasing the number of CPUs on the indexers in the cluster.

Answer: A

Explanation:

Increasing the search factor in the cluster will best address the requirement of high availability for searchable data. The search factor determines how many copies of searchable data are maintained by the cluster. A higher search factor means that more indexers can serve the data in case of a failure or a maintenance event. Increasing the replication factor will improve the availability of raw data, but not searchable data. Increasing the number of search heads or CPUs on the indexers will improve the search performance, but not the availability of searchable data. For more information, see Replication factor and search factor in the Splunk documentation.

3.Search dashboards in the Monitoring Console indicate that the distributed deployment is approaching its capacity.

Which of the following options will provide the most search performance improvement?

- A. Replace the indexer storage to solid state drives (SSD).
- B. Add more search heads and redistribute users based on the search type.
- C. Look for slow searches and reschedule them to run during an off-peak time.
- D. Add more search peers and make sure forwarders distribute data evenly across all indexers.

Answer: D

Explanation:

Adding more search peers and making sure forwarders distribute data evenly across all indexers will provide the most search performance improvement when the distributed deployment is approaching its capacity. Adding more search peers will increase the search concurrency and reduce the load on each

indexer. Distributing data evenly across all indexers will ensure that the search workload is balanced and no indexer becomes a bottleneck. Replacing the indexer storage to SSD will improve the search performance, but it is a costly and time-consuming option. Adding more search heads will not improve the search performance if the indexers are the bottleneck. Rescheduling slow searches to run during an off-peak time will reduce the search contention, but it will not improve the search performance for each individual search. For more information, see [Scale your indexer cluster] and [Distribute data across your indexers] in the Splunk documentation.

4.A Splunk architect has inherited the Splunk deployment at Buttercup Games and end users are complaining that the events are inconsistently formatted for a web source. Further investigation reveals that not all weblogs flow through the same infrastructure: some of the data goes through heavy forwarders and some of the forwarders are managed by another department.

Which of the following items might be the cause of this issue?

- A. The search head may have different configurations than the indexers.
- B. The data inputs are not properly configured across all the forwarders.
- C. The indexers may have different configurations than the heavy forwarders.
- D. The forwarders managed by the other department are an older version than the rest.

Answer: C

Explanation:

The indexers may have different configurations than the heavy forwarders, which might cause the issue of inconsistently formatted events for a web sourcetype. The heavy forwarders perform parsing and indexing on the data before sending it to the indexers. If the indexers have different configurations than the heavy forwarders, such as different props.conf or transforms.conf settings, the data may be parsed or indexed differently on the indexers, resulting in inconsistent events. The search head configurations do not affect the event formatting, as the search head does not parse or index the data. The data inputs configurations on the forwarders do not affect the event formatting, as the data inputs only determine what data to collect and how to monitor it. The forwarder version does not affect the event formatting, as long as the forwarder is compatible with the indexer. For more information, see [Heavy forwarder versus indexer] and [Configure event processing] in the Splunk documentation.

5.A customer has installed a 500GB Enterprise license. They also purchased and installed a 300GB, no enforcement license on the same license master.

How much data can the customer ingest before the search is locked out?

- A. 300GB. After this limit, the search is locked out.
- B. 500GB. After this limit, the search is locked out.
- C. 800GB. After this limit, the search is locked out.
- D. Search is not locked out. Violations are still recorded.

Answer: D

Explanation:

Search is not locked out when a customer has installed a 500GB Enterprise license and a 300GB, no enforcement license on the same license master. The no enforcement license allows the customer to exceed the license quota without locking search, but violations are still recorded. The customer can ingest up to 800GB of data per day without violating the license, but if they ingest more than that, they will incur a violation. However, the violation will not lock search, as the no enforcement license overrides

the enforcement policy of the Enterprise license. For more information, see [No enforcement licenses] and [License violations] in the Splunk documentation.