Question #1 *Topic 1*

How does Monitoring Console (MC) initially identify the server role(s) of a new Splunk Instance?

A. The MC uses a REST endpoint to query the server.

B. Roles are manually assigned within the MC.

C. Roles are read from distsearch.conf.

D. The MC assigns all possible roles by default.

**Suggested Answer:** *C*

*Community vote distribution*

A (100%)

**noysherer** Highly Voted 3 years, 4 months ago

I think its A

upvoted 8 times

**sovip52250** Most Recent 1 year, 10 months ago

Selected Answer: A

A is correct

upvoted 1 times

**Steve2610** 2 years, 1 month ago

Selected Answer: A

Page 56

upvoted 1 times

**Redtonyeah** 2 years, 6 months ago

Selected Answer: A

A is correct one

upvoted 3 times

**SasnycoN** 2 years, 8 months ago

Selected Answer: A

Answer "A" as per page 67 of SCI pdf.

upvoted 3 times

**SasnycoN** 2 years, 8 months ago

Answer A

upvoted 1 times

**carylinlg** 3 years, 4 months ago

I thing D is better

upvoted 1 times

Question #2 *Topic 1*

A customer has asked for a five-node search head cluster (SHC), but does not have the storage budget to use a replication factor greater than 2. They would like to understand what might happen in terms of the users' ability to view historic scheduled search results if they log onto a search head which doesn't contain one of the 2 copies of a given search artifact.
Which of the following statements best describes what would happen in this scenario?

A. The search head that the user has logged onto will proxy the required artifact over to itself from a search head that currently holds a copy. A copy will also be replicated from that search head permanently, so it is available for future use.

B. Because the dispatch folder containing the search results is not present on the search head, the user will not be able to view the search results.

C. The user will not be able to see the results of the search until one of the search heads is restarted, forcing synchronization of all dispatched artifacts across all search heads.

D. The user will not be able to see the results of the search until the Splunk administrator issues the apply shcluster-bundle command on the search head deployer, forcing synchronization of all dispatched artifacts across all search heads.

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**bobixaka** 6 months, 1 week ago

**Selected Answer: A**

Ref:
https://docs.splunk.com/Documentation/Splunk/latest/DistSearch/SHCarchitecture#Artifact_proxying:~:text=In%20addition%2C%20if,number%20of%20cc

upvoted 1 times

**sutcocuk** 1 year, 4 months ago

**Selected Answer: A**

Page 424

upvoted 1 times

**Redtonyeah** 2 years, 6 months ago

**Selected Answer: A**

A is the correct

upvoted 2 times

**SasnycoN** 2 years, 8 months ago

**Selected Answer: A**

Answer "A"

upvoted 2 times

Question #3 *Topic 1*

Monitoring Console (MC) health check configuration items are stored in which configuration file?

A. healthcheck.conf

B. alert_actions.conf

C. distsearch.conf

D. checklist.conf

**Suggested Answer:** *D*

Reference:

https://docs.splunk.com/Documentation/Splunk/8.1.0/DMC/Customizehealthcheck

*Community vote distribution*

D (100%)

**sovip52250** 1 year, 10 months ago

**Selected Answer: D**

D is correct

upvoted 1 times

**Redtonyeah** 2 years, 6 months ago

D is correct

upvoted 3 times

Question #4 *Topic 1*

What should be considered when running the following CLI commands with a goal of accelerating an index cluster migration to new hardware?

```
$SPLUNK_HOME/bin/splunk edit cluster-config -max_peer_build_load 3

$SPLUNK_HOME/bin/splunk edit cluster-config -max_peer_rep_load 6
```

**server.conf**

```
[clustering]

max_peer_build_load = 2

max_peer_rep_load = 5
```

A. Data ingestion rate

B. Network latency and storage IOPS

C. Distance and location

D. SSL data encryption

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**sovip52250** 1 year, 10 months ago

**Selected Answer: B**

B, SCI, always speak about "bandwidth" terms

upvoted 1 times

**Redtonyeah** 2 years, 6 months ago

**Selected Answer: B**

B, SCI, always speak about "bandwidth" terms

upvoted 2 times

Question #5 *Topic 1*

Which statement is true about subsearches?

A. Subsearches are faster than other types of searches.

B. Subsearches work best for joining two large result sets.

C. Subsearches run at the same time as their outer search.

D. Subsearches work best for small result sets.

**Suggested Answer:** *A*

Reference:

https://community.splunk.com/t5/Archive/Looking-for-way-to-explain-why-subsearches-are-so-slow/m-p/479133

*Community vote distribution*

D (100%)

**v12** Highly Voted 3 years, 10 months ago

D is correct.

upvoted 12 times

**bobixaka** Most Recent 6 months, 1 week ago

Selected Answer: D

Ref:

https://docs.splunk.com/Documentation/Splunk/latest/Search/Aboutsubsearches#Subsearch_performance_considerations:~:text=A%20subsearch%20ca

upvoted 1 times

**Simon_UA** 1 year ago

Selected Answer: D

D is correct

upvoted 1 times

**sovip52250** 1 year, 10 months ago

Selected Answer: D

D is correct.

upvoted 1 times

**Steve2610** 2 years, 1 month ago

Selected Answer: D

Page 38

upvoted 1 times

**huu_nguyen** 2 years, 2 months ago

D is correct

upvoted 1 times

**Redtonyeah** 2 years, 6 months ago

Selected Answer: D

D is the correct

upvoted 1 times

**SasnycoN** 2 years, 8 months ago

Answer D is the correct one!

"A" is WRONG!

upvoted 1 times

**splunkingyeti** 3 years, 5 months ago

I would also say that D is correct as well. A is definitely wrong, a sub search is not faster than a tstats search for an example.

upvoted 1 times

**Fake_ID** 3 years, 9 months ago

A seems to be right coz for subsearch

maxtime = <integer>

Maximum number of seconds to run a subsearch before finalizing
Defaults to 60
upvoted 1 times

**jbabbin** 3 years, 10 months ago
A is definitely wrong - Subsearches are notoriously SLOW
See
https://docs.splunk.com/Documentation/Splunk/8.0.2/Search/Aboutsubsearches
Subsearches are mainly used for two purposes:

Parameterize one search, using the output of another search. The example, described above, of searching for the most active host in the last hour is a an example of this use of a subsearch.
Run a separate search and add the output to the first search using the append command.
upvoted 1 times

Question #6 *Topic 1*

A customer has been using Splunk for one year, utilizing a single/all-in-one instance. This single Splunk server is now struggling to cope with the daily ingest rate.
Also, Splunk has become a vital system in day-to-day operations making high availability a consideration for the Splunk service. The customer is unsure how to design the new environment topology in order to provide this.
Which resource would help the customer gather the requirements for their new architecture?

A. Direct the customer to the docs.splunk.com and tell them that all the information to help them select the right design is documented there.

B. Ask the customer to engage with the sales team immediately as they probably need a larger license.

C. Refer the customer to answers.splunk.com as someone else has probably already designed a system that meets their requirements.

D. Refer the customer to the Splunk Validated Architectures document in order to guide them through which approved architectures could meet their requirements.

**Suggested Answer:** *D*
Reference:
https://www.splunk.com/pdfs/technical-briefs/splunk-validated-architectures.pdf

*Community vote distribution*
D (100%)

**spl_bonn** 2 years ago
Selected Answer: D
D is correct.
upvoted 1 times

**Redtonyeah** 2 years, 6 months ago
Selected Answer: D
D, is the correct
upvoted 2 times

**SasnycoN** 2 years, 8 months ago
Confirming D
upvoted 3 times

Question #7 *Topic 1*

The customer has an indexer cluster supporting a wide variety of search needs, including scheduled search, data model acceleration, and summary indexing.
Here is an excerpt from the cluster mater's server.conf:

```
[clustering]
replication_factor=2
search_factor=1
summary_replication-false
```

Which strategy represents the minimum and least disruptive change necessary to protect the searchability of the indexer cluster in case of indexer failure?

A. Enable maintenance mode on the CM to prevent excessive fix-up and bring the failed indexer back online.

B. Leave replication_factor=2, increase search_factor=2 and enable summary_replication.

C. Convert the cluster to multi-site and modify the server.conf to be site_replication_factor=2, site_search_factor=2.

D. Increase replication_factor=3, search_factor=2 to protect the data, and allow there to always be a searchable copy.

**Suggested Answer:** *D*

*Community vote distribution*

B (100%)

**simo988** Highly Voted 3 years, 10 months ago
The correct answer is B
https://docs.splunk.com/Documentation/Splunk/8.1.1/Indexer/Clustersandsummaryreplication
upvoted 9 times

**marinatedcohort** Most Recent 1 month, 2 weeks ago
Selected Answer: D
D - this is the only answer that increases redundancy (replication_factor) which is the main point of the question
upvoted 1 times

**spl_bonn** 2 years ago
Selected Answer: B
I would also say that B is the correct one.
upvoted 2 times

**Redtonyeah** 2 years, 6 months ago
Selected Answer: B
B is the correct
upvoted 2 times

**SasnycoN** 2 years, 8 months ago
Answer is "B" as is "...minimum and least disruptive change"
upvoted 3 times

**jugulinho** 3 years, 6 months ago
D is best choice
upvoted 2 times

**jbabbin** 3 years, 10 months ago
https://docs.splunk.com/Documentation/Splunk/8.0.2/Indexer/Migratetomultisite
upvoted 1 times

**pbandj12** 3 years, 2 months ago
"...minimum and least disruptive change"...migrating to multisite, is not that.
upvoted 1 times

**jbabbin** 3 years, 10 months ago
Splunk reference is for rep factor of 3 in order to tolerate a failure of 2 nodes; which assumes that the cluster is made of of the min of 3 nodes.
B would assume that the cluster is 2 nodes and the nodes would have a search each

https://docs.splunk.com/Documentation/Splunk/8.1.1/Indexer/Thereplicationfactor
upvoted 1 times

Question #8 *Topic 1*

What is the primary driver behind implementing indexer clustering in a customer's environment?

A. To improve resiliency as the search load increases.

B. To reduce indexing latency.

C. To scale out a Splunk environment to offer higher performance capability.

D. To provide higher availability for buckets of data.

**Suggested Answer:** *D*

Reference:

https://docs.splunk.com/Documentation/Splunk/8.1.0/Indexer/Howclusteredsearchworks

*Community vote distribution*

D (100%)

**spl_bonn** 2 years ago

**Selected Answer: D**

D is correct.

upvoted 1 times

**Redtonyeah** 2 years, 6 months ago

**Selected Answer: D**

I would D, base on page 44 SCI

upvoted 4 times

Question #9 *Topic 1*

In a single indexer cluster, where should the Monitoring Console (MC) be installed?

A. Deployer sharing with master cluster.

B. License master that has 50 clients or more.

C. Cluster master node

D. Production Search Head

**Suggested Answer:** *C*
Reference:
https://docs.splunk.com/Documentation/Splunk/8.1.0/DMC/WheretohostDMC

*Community vote distribution*
C (100%)

**spl_bonn** 2 years ago

**Selected Answer: C**

C is fine.

upvoted 1 times

**Redtonyeah** 2 years, 6 months ago

**Selected Answer: C**

C is the correct

upvoted 1 times

**SasnycoN** 2 years, 8 months ago

**Selected Answer: C**

Answer "C"

upvoted 2 times

**iwill_pass** 3 years, 3 months ago

https://docs.splunk.com/Documentation/Splunk/8.1.0/DMC/WheretohostDMC

The master node, if the load on the master node is below the limits specified in Additional roles for the master node in the Managing Indexers and Clusters of Indexers manual. Otherwise, run the monitoring console on a search head node that is dedicated to running monitoring console searches. If you are using SmartStore you must host the monitoring console on a dedicated search head.

upvoted 3 times

Question #10 *Topic 1*

A customer has downloaded the Splunk App for AWS from Splunkbase and installed it in a search head cluster following the instructions using the deployer. A power user modifies a dashboard in the app on one of the search head cluster members. The app containing an updated dashboard is upgraded to the latest version by following the instructions via the deployer.
What happens?

A. The updated dashboard will not be deployed globally to all users, due to the conflict with the power user's modified version of the dashboard.

B. Applying the search head cluster bundle will fail due to the conflict.

C. The updated dashboard will be available to the power user.

D. The updated dashboard will not be available to the power user; they will see their modified version.

**Suggested Answer:** *A*

*Community vote distribution*

D (100%)

**sunil299** Highly Voted 3 years, 3 months ago
i assume D is correct answer. as deployer does not override user/UI modifications
upvoted 5 times

**StevieRayB** Most Recent 1 year, 7 months ago
This is a prime example of reading every word of the question. The "updated" version of the dashboard is referring to the one in the Application update. The one that the power user modified (their) is the one that they will see. The correct answer is "D".
upvoted 1 times

**spl_bonn** 2 years ago
Selected Answer: D
D is the correct one. The local copy on the search head cluster will win.
upvoted 2 times

**RedYeti** 2 years, 2 months ago
Selected Answer: D
Definitely D.
When a user have a role that can modify a dashboard, if he do it, the original xml file is untouched and stay in "default" folder but the new one is in "local" folder and take precedence over the one in "default".
When the app is updated, the original file is updated BUT the file in local is untouched and still have precedence (tested right now on a dev platform).
upvoted 4 times

**Redtonyeah** 2 years, 6 months ago
Selected Answer: D
D, the deployer update default folder,
upvoted 2 times

**Aadil0101010** 2 years, 6 months ago
the correct answer is C
https://docs.splunk.com/Documentation/Splunk/8.2.6/DistSearch/PropagateSHCconfigurationchanges#:~:text=The%20deployer%20is%20a%20Splunk,is%
App upgrades and runtime changes
Because of how configuration file precedence works, changes that users make to apps at runtime get maintained in the apps through subsequent upgrac

Say, for example, that you deploy the 1.0 version of some app, and then a user modifies the app's dashboards. When you later deploy the 1.1 version of t the 1.1 version of the app.
upvoted 3 times

**RedYeti** 2 years, 2 months ago
yes but in that case it is available for every one, not only for power users
upvoted 1 times

Question #11 *Topic 1*

A customer's deployment server is overwhelmed with forwarder connections after adding an additional 1000 clients. The default phone home interval is set to 60 seconds. To reduce the number of connection failures to the DS what is recommended?

A. Create a tiered deployment server topology.

B. Reduce the phone home interval to 6 seconds.

C. Leave the phone home interval at 60 seconds.

D. Increase the phone home interval to 600 seconds.

**Suggested Answer:** *A*

*Community vote distribution*

D (100%)

**IDM** Highly Voted 3 years, 6 months ago

PS recommendation is to always increase the default phone home to 600 sec before leaving site so before adding additional Deployers I would do that..

So D would be my answer

upvoted 12 times

**spl_bonn** Most Recent 2 years ago

Selected Answer: D

D is the correct one.

upvoted 1 times

**pepeperez** 2 years, 4 months ago

Selected Answer: D

D, 600 secs

upvoted 1 times

**Redtonyeah** 2 years, 6 months ago

Selected Answer: D

D, 600 seconds

upvoted 1 times

**SasnycoN** 2 years, 8 months ago

Selected Answer: D

Answer "D"

upvoted 4 times

**jbabbin** 3 years, 10 months ago

D could also be an answer if the customer doesn't have enough resources to create the Tiered Deployment Server Structure.

https://gist.github.com/sworisbreathing/d236895568eaf31da579

IE slowing down the phone home time to 10 minutes would slow down the connection collisions.

Third option not here would be to use DNS name for the DS then utilize Round Robin or some other type of Load Balancing to handle connection requests.

upvoted 4 times

Question #12 *Topic 1*

Which of the following server.conf stanzas indicates the Indexer Discovery feature has not been fully configured (restart pending) on the Master Node?

A.

```
[indexer_discovery]
pass4SymmKey = $7$XcXl1lu3820Jbui14oVe324+mvx6gCKKv6kf2zEaVB6Ie4DcZ647CnLVlfW
```

B.

```
[clustering]
mode = master
pass4SymmKey = $7$tYTXzke+1r+3DULTHHDUTmYOXdtZJPxm21XwMARrJE20jsmicp9C3ni0
```

C.

```
[indexer_discovery]
pass4SymmKey = idxdiscovery
```

D.

```
[clustering]
mode = forwarder
pass4SymmKey = $7$PU9SBXww63Vz3UJdDYGIN0UrdscRh83ssC2pEpwE6P3gn50iNF094g==
```

**Suggested Answer:** *C*

Reference:

https://docs.splunk.com/Documentation/Splunk/8.1.0/Indexer/indexerdiscovery

**pepeperez** Highly Voted 2 years, 4 months ago

C, has not hashed it

upvoted 6 times

**Redtonyeah** Highly Voted 2 years, 6 months ago

C, text is in clear

upvoted 5 times

**spl_bonn** Most Recent 2 years ago

C is correct.

upvoted 1 times

Question #13 *Topic 1*

What is the Splunk PS recommendation when using the deployment server and building deployment apps?

A. Carefully design smaller apps with specific configuration that can be reused.

B. Only deploy Splunk PS base configurations via the deployment server.

C. Use $SPLUNK_HOME/etc/system/local configurations on forwarders and only deploy TAs via the deployment server.

D. Carefully design bigger apps containing multiple configs.

**Suggested Answer:** *B*
Reference:
https://www.splunk.com/en_us/blog/platform/adding-a-deployment-server-forwarder-management-to-a-new-or-existing-splunk-cloud-or-splunk- enterprise-deployment.html

*Community vote distribution*
A (100%)

**sunil299** Highly Voted 3 years, 3 months ago
Answer should be A, as base config can be copied manually to servers too.
upvoted 14 times

**Pete474** Most Recent 1 year, 8 months ago
Selected Answer: A
A is correct
upvoted 1 times

**M9201715** 1 year, 10 months ago
Selected Answer: A
A is definitely correct, apps should always be small and reusable
upvoted 1 times

**spl_bonn** 2 years ago
Selected Answer: A
A is the correct one.
upvoted 1 times

**pepeperez** 2 years, 4 months ago
Selected Answer: A
Better to create specific apps
upvoted 1 times

**Redtonyeah** 2 years, 6 months ago
Selected Answer: A
a is fine
upvoted 3 times

**SasnycoN** 2 years, 8 months ago
I think that's Answer "A"
upvoted 1 times

**jbabbin** 3 years, 10 months ago
Correct link
https://www.splunk.com/en_us/blog/platform/adding-a-deployment-server-forwarder-management-to-a-new-or-existing-splunk-cloud-or-splunk-enterprise-deployment.html
upvoted 1 times

Question #14 *Topic 1*

Which of the following processor occur in the indexing pipeline?

A. tcp out, syslog out

B. Regex replacement, annotator

C. Aggregator

D. UTF-8, linebreaker, header

**Suggested Answer:** *D*
Reference:
https://docs.splunk.com/Documentation/Splunk/8.1.0/Indexer/Howindexingworks#Event_processing_and_the_data_pipeline

*Community vote distribution*
A (100%)

**iwill_pass** Highly Voted 3 years, 3 months ago
Answer is A

- IndexPipe: Tcpout to another Splunk, syslog output, and indexing are done here.
In addition, this pipeline is responsible for bytequota, block signing, and indexing metrics such as thruput etc.
https://wiki.splunk.com/Community:HowIndexingWorks
upvoted 13 times

**marinatedcohort** Most Recent 6 months, 2 weeks ago
Selected Answer: A
A - https://community.splunk.com/t5/Getting-Data-In/Diagrams-of-how-indexing-works-in-the-Splunk-platform-the-Masa/m-p/590774
upvoted 1 times

**marinatedcohort** 6 months, 2 weeks ago
A - https://community.splunk.com/t5/Getting-Data-In/Diagrams-of-how-indexing-works-in-the-Splunk-platform-the-Masa/m-p/590774
upvoted 1 times

**Pete474** 1 year, 8 months ago
Selected Answer: A
A is correct
upvoted 1 times

**spl_bonn** 2 years ago
Selected Answer: A
A is the correct one.
upvoted 1 times

**Steve2610** 2 years, 1 month ago
Page: 12
upvoted 1 times

**RedYeti** 2 years, 2 months ago
Selected Answer: A
Answer is A
Page 243 of Services Core Implementation course
UTF-8, line breaker and header occurs in the parsing pipeline
upvoted 1 times

**Redtonyeah** 2 years, 6 months ago
Selected Answer: A
A, page 243 SCI
upvoted 3 times

**SasnycoN** 2 years, 8 months ago

Selected Answer: A

Answer A

upvoted 1 times

**Nemo72** 3 years, 10 months ago

A is the answer

upvoted 3 times

**Vin1118** 3 years, 10 months ago

A.

https://wiki.splunk.com/Community:HowIndexingWorks

upvoted 2 times

**jbabbin** 3 years, 10 months ago

A - index pipeline tcp out, syslog out , indexer

When D occurs is in the Parsing Pipeline

upvoted 1 times

**simo988** 3 years, 10 months ago

The answer is B.

Source: https://wiki.splunk.com/Community:HowIndexingWorks

upvoted 1 times

**simo988** 3 years, 10 months ago

I meant A.

upvoted 2 times

Question #15 *Topic 1*

Which configuration item should be set to false to significantly improve data ingestion performance?

A. AUTO_KV_JSON

B. BREAK_ONLY_BEFORE_DATE

C. SHOULD_LINEMERGE

D. ANNOTATE_PUNCT

**Suggested Answer:** *C*

Reference:

https://docs.splunk.com/Documentation/Splunk/8.0.6/Data/Configureeventlinebreaking

*Community vote distribution*

C (100%)

**pepeperez** Highly Voted 2 years, 4 months ago

Selected Answer: C

C, to avoid heuristics, and provide precise configs with it

upvoted 5 times

**spl_bonn** Most Recent 2 years ago

Selected Answer: C

C is fine.

upvoted 1 times

**Redtonyeah** 2 years, 6 months ago

Selected Answer: C

C is OK

upvoted 1 times

Question #16 *Topic 1*

A customer has a new set of hardware to replace their aging indexers. What method would reduce the amount of bucket replication operations during the migration process?

A. Disable the indexing ports on the old indexers.

B. Disable replication ports on the old indexers.

C. Put the old indexers into manual detention.

D. Put the old indexers into automatic detention.

**Suggested Answer:** *D*

*Community vote distribution*

C (100%)

**simplekindaman** Highly Voted 3 years, 10 months ago

It's C... can't manually put indexers into auto detention...

upvoted 10 times

**StefanStettin** Most Recent 1 year, 4 months ago

Selected Answer: C

https://docs.splunk.com/Documentation/Splunk/latest/Indexer/Peerdetention?#Manual_detention

upvoted 1 times

**spl_bonn** 2 years ago

Selected Answer: C

C is the correct one.

upvoted 1 times

**RedYeti** 2 years, 2 months ago

Selected Answer: C

Answer is C

upvoted 3 times

**pepeperez** 2 years, 4 months ago

Selected Answer: C

Puting legacy indexers into detention

upvoted 1 times

**SasnycoN** 2 years, 8 months ago

Selected Answer: C

Answer C

upvoted 1 times

**Nemo72** 3 years, 10 months ago

We agree to C

upvoted 2 times

**jbabbin** 3 years, 10 months ago

Splunk official answer is Automatic Detection C

upvoted 2 times

**v12** 3 years, 10 months ago

should be C

upvoted 2 times

Question #17 *Topic 1*

When a bucket rolls from cold to frozen on a clustered indexer, which of the following scenarios occurs?

A. All replicated copies will be rolled to frozen; original copies will remain.

B. Replicated copies of the bucket will remain on all other indexers and the Cluster Master (CM) assigns a new primary bucket.

C. The bucket rolls to frozen on all clustered indexers simultaneously.

D. Nothing. Replicated copies of the bucket will remain on all other indexers until a local retention rule causes it to roll.

**Suggested Answer:** *B*
Reference:
https://docs.splunk.com/Documentation/Splunk/8.1.0/Indexer/Bucketsandclusters

*Community vote distribution*
B (56%) D (44%)

**b5white** 4 months ago
Depends on whether the rolled index is the primary. If yes, then B. But if not, then D.
upvoted 1 times

**bobixaka** 5 months, 4 weeks ago
**Selected Answer: B**
Slide 365 of Core Implementation
I think that a replicated copy can be searchable and non-searchable depending on the SF/RF settings. We assume there is another searchable copy in the cluster. The SF may be = 1, but not very likely.
The last line says: "The only caveat to this is if one of the remaining copies is a searchable copy (and we froze the primary); in this instance, an already-existing searchable copy can be flagged primary (searchable)."
I think it's "B"
upvoted 1 times

**hpbdcb** 9 months, 4 weeks ago
**Selected Answer: B**
B. first of all what happens: "[...]The manager then stops doing fix-ups on that bucket. It operates under the assumption that the other peers will eventually freeze their copies of that bucket as well." but(!!) then it will also do this: "In 6.3 and later, when a primary copy freezes, the cluster reassigns the primary to another searchable copy, if one exists. Searching then continues on that bucket with the new primary copy. When that primary also freezes, the cluster attempts to reassign the primary yet again to another searchable copy. Once all searchable copies of the bucket have been frozen, searching ceases on that bucket." so overall B is the most correct answer here. ref:
https://docs.splunk.com/Documentation/Splunk/9.1.3/Indexer/Bucketsandclusters#How_the_cluster_handles_frozen_buckets
upvoted 1 times

**Simon_UA** 1 year ago
**Selected Answer: D**
Slide 365 of Core Implementation
upvoted 2 times

**srichansen** 1 year, 6 months ago
**Selected Answer: D**
its D. B is a replicated copy and is not searchable.

From the bootcamp slidedeck notes:
Individual hosts manage their own retention, but notify master. May trigger fixup activity. One host may have frozen a bucket (perhaps because of uneven disk usage) but others may still have a live copy. If this other copy is live, then the bucket will remain searchable. If only replica copies are available, the data will no longer be searchable.

Any bucket frozen in a cluster is only frozen *on that instance*, however, the peer will notify the CM. This is so that no fixup activity is done for that bucket any longer. That is, if a bucket is “missing” because it was frozen, it won’t be replaced. The assumption is that the other indexers will freeze it themselves relatively soon. The only caveat to this is if one of the remaining copies is a searchable copy (and we froze the primary); in this instance, an already-existing searchable copy can be flagged primary (searchable).

upvoted 1 times

**bobixaka** 5 months, 4 weeks ago

I think that a replicated copy can be searchable and non-searchable depending on the SF/RF settings. We assume there is another searchable copy in the cluster. The SF may be = 1, but not very likely.
The last line says: "The only caveat to this is if one of the remaining copies is a searchable copy (and we froze the primary); in this instance, an already-existing searchable copy can be flagged primary (searchable)."
I think it's "B"

upvoted 1 times

**KellyPumphrey** 1 year, 8 months ago

**Selected Answer: B**

https://docs.splunk.com/Documentation/Splunk/8.1.0/Indexer/Bucketsandclusters

Note: In 6.3, a change was made in how the cluster responds to frozen primary bucket copies, in order to prolong the time that a bucket remains available for searching:

In a pre-6.3 cluster, when a primary copy freezes, the cluster does not attempt to reassign the primary to any other remaining searchable copy. Searching on a bucket ceases once the primary is frozen.
In 6.3 and later, when a primary copy freezes, the cluster reassigns the primary to another searchable copy, if one exists. Searching then continues on that bucket with the new primary copy. When that primary also freezes, the cluster attempts to reassign the primary yet again to another searchable copy. Once all searchable copies of the bucket have been
frozen, searching ceases on that bucket.

upvoted 3 times

**jcisco123** 1 year, 9 months ago

Answer B: Verified by ChatGPT.

upvoted 2 times

**M9201715** 1 year, 10 months ago

**Selected Answer: B**

I think that B is a better answer. It's basically the same as D, with the addition of the point about primary buckets. And that is exactly what happens if the frozen bucket happens to be primary and there happens to be a replicated bucket that is searchable. The question doesn't specifically say the frozen bucket is primary, but that doesn't mean we should ignore the possibility. Answer B is more complete than D.

upvoted 4 times

**M9201715** 1 year, 10 months ago

Sorry, forgot to add the reference: Services Core Implementation notes, page 365

upvoted 1 times

**spl_bonn** 2 years ago

**Selected Answer: D**

I would say also that D is the correct one. We do not talk about primary bucket.

upvoted 1 times

**Redtonyeah** 2 years, 6 months ago

D is the correct, if in the question speak about primary bucket, would be B, but it is not the case

upvoted 2 times

**SasnycoN** 2 years, 8 months ago

**Selected Answer: D**

Answer D. Source - page 365 of "Services Core Implementation.pdf"

upvoted 3 times

**Vatsal001** 2 years, 10 months ago

D is correct answer.

upvoted 1 times

**iwill_pass** 3 years, 3 months ago

B is correct

when a primary copy freezes, the cluster reassigns the primary to another searchable copy, if one exists. Searching then continues on that bucket with the new primary copy. When that primary also freezes, the cluster attempts to reassign the primary yet again to another searchable copy. Once all searchable copies of the bucket have been frozen, searching ceases on that bucket.

upvoted 4 times

**RedYeti** 2 years, 2 months ago

No, Cluster master just stops doing fixups on the bucket thinking others indexers will eventually freeze their copies too.

upvoted 2 times

Question #18 *Topic 1*

A site from a multi-site indexer cluster needs to be decommissioned. Which of the following actions must be taken?

A. Nothing. Decommissioning a site is not possible.

B. Create an alias for where the new data should be sent.

C. Remove the site from the list of available sites.

D. Remove the site from the list of available sites and create an alias for where the new data should be sent.

**Suggested Answer:** *D*

*Community vote distribution*


**spl_bonn** 2 years ago

Selected Answer: D

D is the correct one.

upvoted 1 times

**leandrojuare2** 1 year, 11 months ago

Hello, could you please explain why D? According to https://docs.splunk.com/Documentation/Splunk/latest/Indexer/Decommissionasite, site mapping does not requiere creating an alias. Thank you.

upvoted 1 times

**bobixaka** 6 months, 1 week ago

I think the site mapping is actually an alias of the site. You map the decommissioned site to an existing one, so forwarders can send the data to the new site, which is aliased ("mapped").

upvoted 1 times

**pepeperez** 2 years, 4 months ago

Selected Answer: D

D is fine

upvoted 3 times

**Redtonyeah** 2 years, 6 months ago

D is OK

upvoted 3 times

**SasnycoN** 2 years, 8 months ago

Selected Answer: D

Answer D

upvoted 4 times

**leandrojuare2** 1 year, 11 months ago

Hello, could you please explain why D? According to https://docs.splunk.com/Documentation/Splunk/latest/Indexer/Decommissionasite, site mapping does not requiere creating an alias. Thank you.

upvoted 1 times

**RedYeti** 1 year, 11 months ago

I wonder if they are confusing alias and map

upvoted 2 times

**carylinlg** 3 years, 6 months ago

site mapping

upvoted 2 times

**jbabbin** 3 years, 10 months ago

Link

https://docs.splunk.com/Documentation/Splunk/8.1.1/Indexer/Decommissionasite

upvoted 1 times

Question #19 *Topic 1*

A customer wants to implement LDAP because managing local Splunk users is becoming too much of an overhead. What configuration details are needed from the customer to implement LDAP authentication?

A. API: Python script with PAM/RADIUS details.

B. LDAP server: port, bind user credentials, path/to/groups, path/to/user.

C. LDAP server: port, bind user credentials, base DN for groups, base DN for users.

D. LDAP REST details, base DN for groups, base DN for users.

**Suggested Answer:** *C*
Reference:
https://www.learnsplunk.com/splunk-ldap-authentication-configuration.html

*Community vote distribution*
C (100%)

**SasnycoN** Highly Voted 2 years, 8 months ago
Selected Answer: C
Answer C
upvoted 7 times

**spl_bonn** Most Recent 2 years ago
Selected Answer: C
C is correct.
upvoted 1 times

**pepeperez** 2 years, 4 months ago
Selected Answer: C
C, all the info is needed to configure LDAP
upvoted 2 times

**Redtonyeah** 2 years, 6 months ago
C is OK
upvoted 2 times

**jbabbin** 3 years, 10 months ago
Splunk Link https://docs.splunk.com/Documentation/Splunk/8.1.1/Security/ConfigureLDAPwithSplunkWeb
upvoted 1 times

Question #20 *Topic 1*

A customer has a search cluster (SHC) of six members split evenly between two data centers (DC). The customer is concerned with network connectivity between the two DCs due to frequent outages. Which of the following is true as it relates to SHC resiliency when a network outage occurs between the two DCs?

A. The SHC will function as expected as the SHC deployer will become the new captain until the network communication is restored.

B. The SHC will stop all scheduled search activity within the SHC.

C. The SHC will function as expected as the minimum required number of nodes for a SHC is 3.

D. The SHC will function as expected as the SHC captain will fall back to previous active captain in the remaining site.

**Suggested Answer:** *D*

*Community vote distribution*

B (100%)

**noysherer** Highly Voted 3 years, 4 months ago

It is B

https://docs.splunk.com/Documentation/Splunk/8.2.1/DistSearch/DeploymultisiteSHC

upvoted 10 times

**sunil299** Highly Voted 3 years, 3 months ago

B seems correct answers, as both site have equal SH and captain cannot be elected.

During this time, the members on the other sites can continue to function as independent search heads. However, they will only be able to service ad hoc searches. Scheduled reports and alerts will not run, because, in a cluster, the scheduling function is relegated to the captain.

upvoted 5 times

**bobixaka** Most Recent 6 months, 1 week ago

Selected Answer: B

To function properly the Captain needs the global majority across all sites to be alive.

When one of the sites fail, we are left with 3 members only, which is not 51% majority and a Captain cannot be elected.

Scheduled searches will be stopped.

Ref:

https://docs.splunk.com/Documentation/Splunk/latest/DistSearch/DeploymultisiteSHC#What_happens_when_the_site_with_the_majority_fails

upvoted 1 times

**spl_bonn** 2 years ago

Selected Answer: B

B is the correct one. SHC will lose the majority to elect a new captain.

upvoted 3 times

**pepeperez** 2 years, 4 months ago

Selected Answer: B

B is the one

upvoted 2 times

**Redtonyeah** 2 years, 6 months ago

I think B is correct

upvoted 3 times

**prich1111** 3 years, 3 months ago

Its C. Both sites will have 3 Shs. Each site will be able to elect a new captain independently and all functionality remains the same

upvoted 4 times

**k3115807** 3 years ago

I found that. shc in 6 members need 4 members to election.

https://docs.splunk.com/Documentation/Splunk/8.2.1/DistSearch/SHCarchitecture#Captain_election

To become captain, a member needs to win a majority vote of all members. For example, in a seven-member cluster, election requires four votes. Similarly, a six-member cluster also requires four votes.

upvoted 1 times

Question #21 *Topic 1*

A [script://] input sends data to a Splunk forwarder using which method?

A. UDP stream

B. TCP stream

C. Temporary file

D. STDOUT/STDERR

**Suggested Answer:** *C*

Reference:

https://docs.splunk.com/Documentation/Splunk/latest/Admin/inputsconf

*Community vote distribution*

D (100%)

**jbabbin** Highly Voted 3 years, 10 months ago

Looks like D is correct since the script input passes results/output to STDOUT for ingestion to splunk

https://docs.splunk.com/Documentation/Splunk/latest/AdvancedDev/ScriptWriting

upvoted 7 times

**aymenbest2** Most Recent 1 year, 6 months ago

Selected Answer: D

Splunk runs the script [script://] stanza and ingest the outpu (STDOUT/STDERR)

upvoted 1 times

**spl_bonn** 2 years ago

D is the correct one. Page 97 of Splunk Core Implementation

upvoted 1 times

**pepeperez** 2 years, 4 months ago

Selected Answer: D

D, script outputs on that format

upvoted 2 times

**Redtonyeah** 2 years, 6 months ago

D is the correct

upvoted 1 times

**chuchoneitor** 3 years, 2 months ago

The correct one is D.

upvoted 2 times

**jbabbin** 3 years, 10 months ago

C is correct script write to a Temp File as most scripted inputs write to a file that then is indexed

a script would go to STDOUT/STDERR only if it was a HEC like connection or Splunk RESTful input

upvoted 1 times

**v12** 3 years, 10 months ago

D looks correct

upvoted 3 times

Question #22 *Topic 1*

A customer wants to understand how Splunk bucket types (hot, warm, cold) impact search performance within their environment. Their indexers have a single storage device for all data. What is the proper message to communicate to the customer?

A. The bucket types (hot, warm, or cold) have the same search performance characteristics within the customer's environment.

B. While hot, warm, and cold buckets have the same search performance characteristics within the customers environment, due to their optimized structure, the thawed buckets are the most performant.

C. Searching hot and warm buckets result in best performance because by default the cold buckets are miniaturized by removing TSIDX files to save on storage cost.

D. Because the cold buckets are written to a cheaper/slower storage volume, they will be slower to search compared to hot and warm buckets which are written to Solid State Disk (SSD).

**Suggested Answer:** *D*

*Community vote distribution*

A (100%)

**v12** Highly Voted 3 years, 10 months ago

A looks correct

upvoted 8 times

**736b2ff** Most Recent 3 months, 3 weeks ago

Selected Answer: A

let's fix this

upvoted 1 times

**hpbdcb** 9 months, 4 weeks ago

Selected Answer: A

A because TSIDX removal only happens when configured (https://docs.splunk.com/Documentation/Splunk/9.1.3/Indexer/Reducetsidxdiskusage) and as bucket states are stored on the same storage A is the only correct answer here.

upvoted 1 times

**spl_bonn** 2 years ago

A is the correct one.

upvoted 1 times

**Swagdadp215** 2 years, 5 months ago

Selected Answer: A

A is correct. Since the buckets will reside of the same storage device, they should have the same performance characteristics. The others are irrelevant due to that input from the question.

upvoted 2 times

**RedYeti** 2 years, 2 months ago

What about the C?

upvoted 1 times

**Redtonyeah** 2 years, 6 months ago

A is OK

upvoted 1 times

**jbabbin** 3 years, 10 months ago

Agreed!

D has no context and assumes not in the question that the customer has different storage but the question says it's all on ONE hardware platform

upvoted 1 times

**simplekindaman** 3 years, 10 months ago

I agree... if they have all the data on the same type of disk, performance will be the same for bucket states... it is A

upvoted 1 times

Question #23 *Topic 1*

An index receives approximately 50GB of data per day per indexer at an even and consistent rate. The customer would like to keep this data searchable for a minimum of 30 days. In addition, they have hourly scheduled searches that process a week's worth of data and are quite sensitive to search performance.
Given ideal conditions (no restarts, nor drops/bursts in data volume), and following PS best practices, which of the following sets of indexes.conf settings can be leveraged to meet the requirements?

A. frozenTimePeriodInSecs, maxDataSize, maxVolumeDataSizeMB, maxHotBuckets

B. maxDataSize, maxTotalDataSizeMB, maxHotBuckets, maxGlobalDataSizeMB

C. maxDataSize, frozenTimePeriodInSecs, maxVolumeDataSizeMB

D. frozenTimePeriodInSecs, maxWarmDBCount, homePath.maxDataSizeMB, maxHotSpanSecs

**Suggested Answer:** *B*

*Community vote distribution*


**bobixaka** 6 months, 1 week ago

**Selected Answer: C**

In the Splunk PS Base App org_all_indexes in indexes.conf you have this:
[customer_index]
disabled = false
homePath = volume:primary/$_index_name/db
coldPath = volume:primary/$_index_name/colddb
thawedPath = volume:primary/$_index_name/thaweddb
frozenTimePeriodInSecs = 31536000
maxTotalDataSizeMB = 102400
maxDataSize = auto/auto_high_volume

That's all you need. As mentioned everything works fine, no drops, no restarts no bursts of data. You have ideal conditions, so you can use all default settings.
50GB/day is not so much for an indexer to handle using default settings.
You don't need to specify maxHotBuckets, the default "auto" value is fine here. A is not correct.
You don't need maxGlobalDataSizeMB here, because it's for SmartStore. B is not correct.
You don't need to specify maxWarmDBCount and maxHotSpanSecs, so D is not correct.
upvoted 1 times

**BMS0598** 1 year, 8 months ago

**Selected Answer: B**

Considering B as the correct answer because the question states that the data should be searchable for a MINIMUM of 30 days, but it doesn't necessarily mean the buckets older than 30 days need to be moved to frozen.
Option B focuses more on settings related to storage which could be used to set an appropriate storage size limit if you do the math correctly based on the fact they tell you that 50GB are ingested per day per indexer and that no drops/bursts of data will happen. So the question seems to be more focused on using storage-related settings to keep the buckets searchable for at least 30 days.
Also, chatGPT suggests the correct answer is B.
upvoted 1 times

**spl_bonn** 2 years ago

**Selected Answer: A**

I agree also that A is the correct one.
upvoted 2 times

**RedYeti** 2 years, 2 months ago

**Selected Answer: A**

maxDataSize:
Maximum size that a hot bucket can reach before rolling to warm, use auto_high_volume for high-volume indexes (>10GB per day).

maxHotBuckets:
maximum number of hot buckets that can exist per index.

frozenTimePeriodInSecs:
number of seconds after which indexed data rolls to frozen.

---

"50GB of data per day per indexer" means to set maxDataSize to auto_high_volume and maxHotBuckets not to auto but to a high value (10)
=> not answer D and C

"data searchable for a minimum of 30 days" means to set frozenTimePeriodInSecs
=> not answer B

So answer is A
upvoted 3 times

**Redtonyeah** 2 years, 6 months ago
A is more completed than C, so A
upvoted 2 times

**LearningDani** 2 years, 9 months ago
I think it's A

> 'maxDataSize' should be to 'auto_high_volume', because A "high volume index" would typically be considered one that gets over 10GB of data per day.
> 'maxHotBuckets' - this defines the maximum number of simultaneously open hot buckets (actively being written to). For indexes that receive a lot of data, this should be 10, other indexes can safely keep the default
> 'maxVolumeDataSizeMB' should also be set to define the size the volume
> 'frozenTimePeriodInSecs' for sure to set retention time
upvoted 3 times

**sunil299** 3 years, 3 months ago
maxGlobalDataSizeMB is primarily for smartstore.
upvoted 2 times

**noysherer** 3 years, 5 months ago
I think it is D
https://docs.splunk.com/Documentation/Splunk/8.2.0/Admin/Indexesconf
upvoted 3 times

Question #24 *Topic 1*

A customer has a Universal Forwarder (UF) with an inputs.conf monitoring its splunkd.log. The data is sent through a heavy forwarder to an indexer.
Where does the Index time parsing occur?

A. Indexer

B. Universal forwarder

C. Search head

D. Heavy forwarder

**Suggested Answer:** *D*
Reference:
https://www.learnsplunk.com/splunk-interview-questions.html

*Community vote distribution*
D (100%)

**pepeperez** Highly Voted 2 years, 4 months ago
Selected Answer: D
Heavy forwarder, does parsing before idx
upvoted 5 times

**spl_bonn** Most Recent 2 years ago
Selected Answer: D
D is the correct one. Parsing takes place on the first full Splunk instance.
upvoted 1 times

**pentel** 2 years, 5 months ago
Selected Answer: D
D, in the HF
upvoted 2 times

**Redtonyeah** 2 years, 6 months ago
A, in the IDX
upvoted 1 times

**Redtonyeah** 2 years, 6 months ago
D, in the HF
upvoted 2 times

Question #25 *Topic 1*

The customer wants to migrate their current Splunk Index cluster to new hardware to improve indexing and search performance. What is the correct process and procedure for this task?

A. 1. Install new indexers. 2. Configure indexers into the cluster as peers; ensure they receive the same configuration via the deployment server. 3. Decommission old peers one at a time. 4. Remove old peers from the CM's list. 5. Update forwarders to forward to the new peers.

B. 1. Install new indexers. 2. Configure indexers into the cluster as peers; ensure they receive the cluster bundle and the same configuration as original peers. 3. Decommission old peers one at a time. 4. Remove old peers from the CM's list. 5. Update forwarders to forward to the new peers.

C. 1. Install new indexers. 2. Configure indexers into the cluster as peers; ensure they receive the same configuration via the deployment server. 3. Update forwarders to forward to the new peers. 4. Decommission old peers on at a time. 5. Restart the cluster master (CM).

D. 1. Install new indexers. 2. Configure indexers into the cluster as peers; ensure they receive the cluster bundle and the same configuration as original peers. 3. Update forwarders to forward to the new peers. 4. Decommission old peers one at a time. 5. Remove old peers from the CM's list.

**Suggested Answer:** *C*

*Community vote distribution*


**Vin1118** Highly Voted 3 years, 10 months ago

Should be D.

Clustered Indexers will not received configs from DS

upvoted 12 times

**spl_bonn** Most Recent 2 years ago

Selected Answer: D

It should be D.

upvoted 2 times

**pepeperez** 2 years, 4 months ago

Selected Answer: D

D is the correct order

upvoted 3 times

**Redtonyeah** 2 years, 6 months ago

D is fine

upvoted 1 times

**LearningDani** 2 years, 9 months ago

I also go with D

I think the Forwarders should be updated befor decomissioning the old peers so that they do not try to send to already decomissioned systems...

upvoted 2 times

**simplekindaman** 3 years, 10 months ago

Agree that A and C are not valid because of the DS. I will tend to agree with D, but B is valid also.

upvoted 1 times

**jbabbin** 3 years, 10 months ago

C and A are wrong as the DS is not used for Index Cluster management or IDX apps

B could also be be used if want to age out data on the existing peers or ensure that the new peers are receiving data before taking the existing peers offline

upvoted 2 times

Question #26 *Topic 1*

Consider the scenario where the /var/log directory contains the files secure, messages, cron, audit. A customer has created the following inputs.conf stanzas in the same Splunk app in order to attempt to monitor the files secure and messages:

```
[monitor:///var/log]
sourcetype = syslog
index = secrutiy
disabled = false
whitelist = messages

[monitor:///var/log]
sourcetype  = syslog
index = security
disabled = false
whitelist = secure
```

Which file(s) will actually be actively monitored?

A. /var/log/secure

B. /var/log/messages

C. /var/log/messages, /var/log/cron, /var/log/audit, /var/log/secure

D. /var/log/secure, /var/log/messages

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**v12** Highly Voted 3 years, 10 months ago

if stanzas are same only the last one gets applied, see the discussions here:-
https://community.splunk.com/t5/Archive/Multiple-stanza-in-inputs-conf-for-the-same-folder/m-p/353748

upvoted 7 times

**simplekindaman** Highly Voted 3 years, 10 months ago

Agree with v12 on this one. The second stanza will override the first, and only secure will be monitored. A is correct

upvoted 7 times

**spl_bonn** Most Recent 2 years ago

Selected Answer: A

A is correct

upvoted 1 times

**pepeperez** 2 years, 4 months ago

Selected Answer: A

page 193 of SCI

upvoted 3 times

**Redtonyeah** 2 years, 6 months ago

A is the correct

upvoted 1 times

**jbabbin** 3 years, 10 months ago

Also both files will be monitored , though the first stanza won't log to splunk assuming the spelling issue with the index, but will be monitored and just have the data lost/not written.

upvoted 1 times

**jbabbin** 3 years, 10 months ago

This is wrong the correct answer is D both of the files would be indexed
Assuming the spelling error with the first stanza is fixed
the whitelist option specifically calls out both files
https://docs.splunk.com/Documentation/Splunk/8.1.1/Data/Whitelistorblacklistspecificincomingdata

upvoted 2 times

Question #27 *Topic 1*

A customer has written the following search:

```
sourcetype=purchase:orders
| table _time, customer, product, amount, order_id
| stats count sum(amount) AS amount latest (_time) AS _time by customer, order_id
| search customer= "timmy*"
| lookup vip_customers customer OUTPUT vip_status
| table _time, customer, order_id, amount, vip_status
| search vip_status= "true"
```

How can the search be rewritten to maximize efficiency?

A.

```
index=sales sourcetype=purchase:orders
| table _time, customer, product, amount, order_id
| stats count sum(amount) AS amount latest (_time) AS _time by customer, order_id
| search customer= "timmy*"
| lookup vip_customers customer OUTPUT vip_status
| table _time, customer, order_id, amount, vip_status
| search vip_status= "true"
```

B.

```
index=proxy source=proxy:data:syslog user= "timmy*"
| table _time, user, url, duration, category, action
| stats count sum(duration) AS duration last(url) AS url latest (_time) AS _time by user
| lookup user_status user OUTPUT status
| table _time, user, status
```

C.

```
index=sales sourcetype=purchase:orders customer= "timmy*"
| lookup vip_customers customer OUTPUT vip_status
| search vip_status= "true"
| stats sum(amount) AS amount latest (_time) AS _time by customer, order_id
| table _time, customer, order_id, amount
```

D.

```
index=sales sourcetype=purchase:orders customer= "timmy*"
| lookup vip_customers customer OUTPUT vip_status
| stats count sum(amount) AS amount latest (_time) AS _time by customer, order_id
| search vip_status= "true"
| table _time, customer, order_id, amount, vip_status
```

**Suggested Answer:** *C*

**hpbdcb** 9 months, 4 weeks ago

must be C while it will NOT result in the same table (missing vip_status field). it must be "customer=" in the main search to limit and D won't work as vip_status is not in the stats command

upvoted 1 times

**spl_bonn** 2 years ago

C is correct.

upvoted 1 times

**Redtonyeah** 2 years, 6 months ago

C is right, the filter always first,

upvoted 4 times

**SasnycoN** 2 years, 8 months ago

Correct answer is "D"

upvoted 1 times

**saraque** 11 months, 3 weeks ago

Nop, it's C. The stats command is not defining the vip_customer field. In that case you will not see results because the search command is looking for a inexistent field.

upvoted 2 times

**jbabbin** 3 years, 10 months ago

Wrong forgot to put the index of sales in the question

upvoted 1 times

**hpbdcb** 9 months, 4 weeks ago

u know about default searched indexes if no index is specified?

upvoted 1 times

Question #28 *Topic 1*

How could a role in which all users must specify an index=clause in all searches be configured?

A. Set the authorize.conf setting: srchIndexesDefault to no value.

B. Set the authorize.conf setting: srchFilter to no value.

C. Set the authorize.conf setting: srchIndexesAllowed to no value.

D. Set the authorize.conf setting: srchJobsQuota to no value.

**Suggested Answer:** *B*

*Community vote distribution*

A (100%)

**jbabbin** Highly Voted 3 years, 10 months ago

Should be A

https://community.splunk.com/t5/Archive/srchIndexesDefault-parameter-is-not-respected-when-srchFilter-is/m-p/495869

The B option only limits what indexes they search by default

upvoted 15 times

**spl_bonn** Most Recent 2 years ago

The correct answer is A.

upvoted 1 times

**pepeperez** 2 years, 4 months ago

Selected Answer: A

Correct is A, if you dont give default indexes

upvoted 3 times

**Redtonyeah** 2 years, 6 months ago

A is the correct

upvoted 2 times

**SasnycoN** 2 years, 8 months ago

Selected Answer: A

I think that the answer is "A"

upvoted 2 times

Question #29 *Topic 1*

In which of the following scenarios should base configurations be used to provide consistent, repeatable, and supportable configurations?

A. For non-production environments to keep their configurations in sync.

B. To ensure every customer has exactly the same base settings.

C. To provide settings that do not need to be customized to meet customer requirements.

D. To provide settings that can be customized to meet customer requirements.

**Suggested Answer:** *C*

Reference:

https://docs.splunk.com/Documentation/Splunk/latest/Admin/Wheretofindtheconfigurationfiles

*Community vote distribution*

D (100%)

**hpbdcb** 9 months, 4 weeks ago

**Selected Answer: D**

each customer has different requirements and doing base configurations depends on them. So D

upvoted 1 times

**mjtitan** 2 years, 2 months ago

I think D

upvoted 1 times

**huu_nguyen** 2 years, 2 months ago

D is the one

upvoted 1 times

**pepeperez** 2 years, 4 months ago

**Selected Answer: D**

I would say D,

upvoted 3 times

**SasnycoN** 2 years, 8 months ago

I think that the answer is "D"

upvoted 4 times

Question #30 *Topic 1*

Data can be onboarded using apps, Splunk Web, or the CLI.
Which is the PS preferred method?

A. Create UDP input port 9997 on a UF.

B. Use the add data wizard in Splunk Web.

C. Use the inputs.conf file.

D. Use a scripted input to monitor a log file.

**Suggested Answer:** *B*
Reference:
https://docs.splunk.com/Documentation/Splunk/8.1.0/Data/Howdoyouwanttoadddata

*Community vote distribution*
C (100%)

**tico2k** Highly Voted 3 years, 10 months ago
PS recommends C
upvoted 10 times

**jbabbin** Highly Voted 3 years, 10 months ago
wrong the Data Wizard is for trialing small sets of data <500MB of log files
The proper method is to use inputs.conf in combination with small single use Splunk apps for the data source.
upvoted 5 times

**spl_bonn** Most Recent 2 years ago
Selected Answer: C
C is the recommended way for Splunk PS
upvoted 2 times

**pepeperez** 2 years, 4 months ago
Selected Answer: C
ALWAYS USE INPUTS! C
upvoted 1 times

**Redtonyeah** 2 years, 6 months ago
C looks fine
upvoted 1 times

**splunkingyeti** 3 years, 5 months ago
Agree that it should be C.
upvoted 4 times

**Vin1118** 3 years, 10 months ago
Should be C. PS recommends base configs apps
upvoted 5 times

Question #31 *Topic 1*

Which of the following statements applies to indexer discovery?

A. The Cluster Master (CM) can automatically discover new indexers added to the cluster.

B. Forwarders can automatically discover new indexers added to the cluster.

C. Deployment servers can automatically configure new indexers added to the cluster.

D. Search heads can automatically discover new indexers added to the cluster.

**Suggested Answer:** *D*
Reference:
https://docs.splunk.com/Documentation/Splunk/8.1.0/DistSearch/Connectclustersearchheadstosearchpeers

*Community vote distribution*
B (100%)

**Nemo72** Highly Voted 3 years, 10 months ago
B is the answer
upvoted 7 times

**hpbdcb** Most Recent 9 months, 4 weeks ago
Selected Answer: B
B (https://docs.splunk.com/Documentation/Splunk/9.2.0/Indexer/indexerdiscovery) and btw not D as the SH will use the Cluster Leader to detect available peers
upvoted 1 times

**spl_bonn** 2 years ago
B is the correct one.
upvoted 1 times

**mjtitan** 2 years, 2 months ago
Selected Answer: B
its B.
upvoted 2 times

**pepeperez** 2 years, 4 months ago
Selected Answer: B
B is thew one
upvoted 2 times

**Redtonyeah** 2 years, 6 months ago
B is OK
upvoted 1 times

**tico2k** 3 years, 10 months ago
The correct answer is B
upvoted 2 times

**jbabbin** 3 years, 10 months ago
B
more info

Briefly, the process works like this:

1. The peer nodes provide the manager node with information on their receiving ports.

2. The forwarders poll the manager at regular intervals for the list of available peer nodes.
upvoted 3 times

**simo988** 3 years, 10 months ago

The answer is B.

Reference: https://docs.splunk.com/Documentation/Splunk/8.1.1/Indexer/indexerdiscovery

upvoted 3 times

Question #32 *Topic 1*

The data in Splunk is now subject to auditing and compliance controls. A customer would like to ensure that at least one year of logs are retained for both
Windows and Firewall events. What data retention controls must be configured?

A. maxTotalDataSizeMB and frozenTimePeriodInSecs

B. coldToFrozenDir and coldToFrozenScript

C. Splunk Volume and maxTotalDataSizMB

D. Splunk Volume and frozenTimePeriodInSecs

**Suggested Answer:** *A*
Reference:
https://docs.splunk.com/Documentation/Splunk/8.1.0/Indexer/Setaretirementandarchivingpolicy

*Community vote distribution*

A (100%)

**RedYeti** 1 year, 11 months ago
'e' is missing in "maxTotalDataSizMB"
upvoted 1 times

**spl_bonn** 2 years ago
**Selected Answer: A**
A is correct.
upvoted 1 times

**huu_nguyen** 2 years, 2 months ago
**Selected Answer: A**
A is the one
upvoted 3 times

**pepeperez** 2 years, 4 months ago
**Selected Answer: A**
I would say A, Splunk volume is nothing specific
upvoted 2 times

**Redtonyeah** 2 years, 6 months ago
A, the rest of configuration doesnt have sense
upvoted 1 times

Question #33 *Topic 1*

What happens when an index cluster peer freezes a bucket?

A. All indexers with a copy of the bucket will delete it.

B. The cluster master will ensure another copy of the bucket is made on the other peers to meet the replication settings.

C. The cluster master will no longer perform fix-up activities for the bucket.

D. All indexers with a copy of the bucket will immediately roll it to frozen.

**Suggested Answer:** *C*

Reference:

https://docs.splunk.com/Documentation/Splunk/8.1.0/Indexer/Bucketsandclusters

*Community vote distribution*

C (100%)

**spl_bonn** 2 years ago

C is correct

upvoted 2 times

**ranga19** 2 years, 2 months ago

C , The cluster master will no longer perform fix-up activities for the bucket.

upvoted 2 times

**huu_nguyen** 2 years, 2 months ago

**Selected Answer: C**

C is correct

upvoted 2 times

**Redtonyeah** 2 years, 6 months ago

C is OK

upvoted 1 times

Question #34 *Topic 1*

A customer has the following Splunk instances within their environment: An indexer cluster consisting of a cluster master/master node and five clustered indexers, two search heads (no search head clustering), a deployment server, and a license master. The deployment server and license master are running on their own single-purpose instances. The customer would like to start using the Monitoring Console (MC) to monitor the whole environment.
On the MC instance, which instances will need to be configured as distributed search peers by specifying them via the UI using the settings menu?

A. Just the cluster master/master node.

B. Indexers, search heads, deployment server, license master, cluster master/master node.

C. Search heads, deployment server, license master, cluster master/master node

D. Deployment server, license master

**Suggested Answer:** *C*

*Community vote distribution*

C (100%)

**ranga19** Highly Voted 2 years, 2 months ago

C, except indexers all other nodes must be configured as distsearch peers because after adding CM /MN all the indexers will be added to MC by default.

upvoted 5 times

**spl_bonn** Most Recent 2 years ago

Selected Answer: C

C is the correct answer

upvoted 1 times

**pepeperez** 2 years, 4 months ago

Selected Answer: C

C includes all instances needed

upvoted 1 times

**Redtonyeah** 2 years, 6 months ago

C is the correct

upvoted 1 times

**jbabbin** 3 years, 10 months ago

Link https://docs.splunk.com/Documentation/Splunk/8.1.1/DMC/Addinstancesassearchpeers

upvoted 3 times

Question #35 *Topic 1*

What does Splunk do when it indexes events?

A. Extracts the top 10 fields.

B. Extracts metadata fields such as host, source, sourcetype.

C. Performs parsing, merging, and typing processes on universal forwarders.

D. Create report acceleration summaries.

**Suggested Answer:** *B*

Reference:

https://docs.splunk.com/Documentation/Splunk/8.1.0/Indexer/Howindexingworks#:~:text=Splunk%20Enterprise%20can%20index%20any,events%

20indexes%20and%20metrics%20indexes

*Community vote distribution*

B (100%)

**spl_bonn** 2 years ago

**Selected Answer: B**

B is correct

upvoted 3 times

**ranga19** 2 years, 2 months ago

B, Extracts metadata fields such as host, source, sourcetype from the Data.

upvoted 2 times

**Redtonyeah** 2 years, 6 months ago

B is the correct

upvoted 1 times

Question #36 *Topic 1*

What is the default push mode for a search head cluster deployer app configuration bundle?

A. full

B. merge_to_default

C. default_only

D. local_only

**Suggested Answer:** *B*

Reference:
https://docs.splunk.com/Documentation/Splunk/8.1.0/DistSearch/PropagateSHCconfigurationchanges#:~:text=The%20deployer%20push%20mode%
20determines,default%20push%20mode%20is%20merge_to_default%20

*Community vote distribution*

B (100%)

**spl_bonn** 2 years ago

**Selected Answer: B**

B is the correct one.

upvoted 3 times

**ranga19** 2 years, 2 months ago

B , merge_to_default on SHC members.

upvoted 2 times

**pepeperez** 2 years, 4 months ago

**Selected Answer: B**

B is the one

upvoted 2 times

**Redtonyeah** 2 years, 6 months ago

B, by default merge all in de default directory

upvoted 1 times

Question #37 *Topic 1*

In which of the following scenarios is a subsearch the most appropriate?

A. When joining results from multiple indexes.

B. When dynamically filtering hosts.

C. When filtering indexed fields.

D. When joining multiple large datasets.

**Suggested Answer:** *A*

*Community vote distribution*

B (100%)

**chuchoneitor** Highly Voted 3 years, 2 months ago

I think B is better because:

• Used to produce search terms for the outer search
– Find a subset of hosts
– Programmatically determine “earliest” and “latest”
– Craft the main search string dynamically
– Subsearches always run first, before the main search

upvoted 11 times

**spl_bonn** Most Recent 2 years ago

B is the correct answer

upvoted 1 times

**mjtitan** 2 years, 2 months ago

Selected Answer: B

B is correct

A - You don't need subsearch to combine multiple indexes

subsearches are not for larger result sets

upvoted 3 times

**Redtonyeah** 2 years, 6 months ago

B, page 320 SCI

upvoted 3 times

**SasnycoN** 2 years, 8 months ago

Selected Answer: B

Also think that the answer is "B"

upvoted 1 times

Question #38 *Topic 1*

A customer has implemented their own Role Based Access Control (RBAC) model to attempt to give the Security team different data access than the Operations team by creating two new Splunk roles "" security and operations. In the srchIndexesAllowed setting of authorize.conf, they specified the network index under the security role and the operations index under the operations role. The new roles are set up to inherit the default user role.
If a new user is created and assigned to the operations role only, which indexes will the user have access to search?

A. operations, network, _internal, _audit

B. operations

C. No Indexes

D. operations, network

**Suggested Answer:** *A*

*Community vote distribution*


**Vin1118** Highly Voted 3 years, 10 months ago
I think D. By default, user have access to all non-internal indexes
upvoted 10 times

**da_stingo** Most Recent 1 year, 4 months ago
Selected Answer: B
I am really struggling with your answers, boys. When I look up the users role which comes preinstalled with every splunk install by default, I do NOT see that users have access to all non-indexes. Do you have another source of truth for answering this question beside looking of the role info page showing which index the role is allowed to search?
upvoted 1 times

**da_stingo** 1 year, 4 months ago
*non-internal-indexes
upvoted 1 times

**da_stingo** 1 year, 4 months ago
never mind. i found it :/
upvoted 1 times

**spl_bonn** 2 years ago
Selected Answer: D
D is the correct one. user role is by default has access to all non-internal indexes
upvoted 1 times

**ranga19** 2 years, 2 months ago
D ,By default, user have access to all non-internal indexes.
upvoted 1 times

**huu_nguyen** 2 years, 2 months ago
Selected Answer: D
D is the one, by default, the user does have access to all non-internal indexes but not to internal indexes
upvoted 1 times

**Redtonyeah** 2 years, 6 months ago
D is OK
upvoted 1 times

**nutsu** 3 years, 1 month ago
answer D,
default user role not see internal index and default see non-internal index
upvoted 2 times

**simplekindaman** 3 years, 10 months ago

Agreed... answer is D, for reason Vin1118 mentions

upvoted 1 times

**tico2k** 3 years, 10 months ago

The correct answer is D

upvoted 1 times

Question #39 *Topic 1*

A customer would like Splunk to delete files after they've been ingested. The Universal Forwarder has read/write access to the directory structure. Which input type would be most appropriate to use in order to ensure files are ingested and then deleted afterwards?

A. Script

B. Batch

C. Monitor

D. Fschange

**Suggested Answer:** *B*

Reference:

https://community.splunk.com/t5/Getting-Data-In/Is-it-possible-to-have-a-Splunk-universal-forwarder-read-a/td-p/172752

*Community vote distribution*

B (100%)

**ranga19** 2 years, 2 months ago

Batch will delete the file after ingesting

upvoted 3 times

**pepeperez** 2 years, 4 months ago

**Selected Answer: B**

Batch will delete it after ingesting

upvoted 2 times

**Redtonyeah** 2 years, 6 months ago

B is the correct

upvoted 1 times

**Redtonyeah** 2 years, 6 months ago

B, after read removed the files

upvoted 2 times