Question #1 *Topic 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.

Suggested Answer: *D*
Reference:
https://docs.splunk.com/Documentation/Splunk/7.3.2/Indexer/Systemrequirements

Community vote distribution
A (94%) 6%

Crash_Override Highly Voted 2 years, 9 months ago

Selected Answer: A

Replicated copies of non-searchable data are smaller than copies of searchable data, because they include only the data and not the associated index files. So setting search factor to n-1 has a greater reduction.

upvoted 12 times

62d8e4c Most Recent 3 months ago

Selected Answer: D

If you only have one copy (RF = 1) by default your SF is 1, so the right answer is D.

upvoted 1 times

bobixaka 9 months, 4 weeks ago

Selected Answer: A

TSIDX files are bigger than the raw data, so A is the correct answer.
If you reduce the number of searchable data copies you will have a greater impact on storage savings. We are assuming that the current search_factor=N and replication_factor=N.

upvoted 1 times

FreshLearn 10 months ago

IMHO the question is tricky and has a massive 'it depends' in it, but considering N could be any cluster size,
let's assume we have a requirement of 2 searchable copies and a replication factor of 4(=4x raw data distributed accross the nodes total) and a cluster with 50 peer nodes => so you set it with
A) to 49(!) full sized searchable copies
D) to 49(!) smaller sized raw copies
in contrast "C) Decreasing the data model acceleration range." would actually DECREASE consumed storage

upvoted 1 times

_bert 1 year, 3 months ago

Replication factor (RF) is the number of copies of a bucket. Search factor (SF) is the number of those copies which are searchable. You can't search more copies than you have so SF must be less than or equal to RF. By reducing RF to N-1 you are automatically reducing the SF to N-1. So answer is D.

upvoted 2 times

Sledge_Hammer 1 year, 6 months ago

D is the right answer. Replication factor is RECOMMENDED to be set at 1 less than the number of peer nodes (N-1)

upvoted 1 times

Harllen91 1 year, 5 months ago

Architecting Splunk Enterprise Deployments, "Best Practice: Minimum (RF +1) peer nodes. You have a replication factor of 3, then you want 4 peer nodes, not one less.

upvoted 1 times

stwong 2 years ago

Seems the question assumes SF is always smaller than RF, while RF in example in
https://docs.splunk.com/Documentation/Splunk/7.3.2/Indexer/Systemrequirements equal to cluster size, e.g.
3 peer nodes, with replication factor = 3; search factor = 2
5 peer nodes, with replication factor = 5; search factor = 3

In these cases, setting SF=N-1 will increase disk usage instead, while since RF=N in both cases, setting RF=N-1 will reduce disk usage.

Is that why having answer = D ?

upvoted 3 times

javo_dlg 2 years, 2 months ago

Answer A is correct

upvoted 2 times

SplunkStreamer 2 years, 3 months ago

Selected Answer: A

Answer: A

upvoted 1 times

frappe 2 years, 4 months ago

Selected Answer: A

Answer A seems correct as Replicated copies of non-searchable data are smaller than copies of searchable data, because they include only the data and not the associated index files.

upvoted 1 times

RedYeti 2 years, 7 months ago

Selected Answer: A

Answer A

upvoted 2 times

manu78 3 years, 8 months ago

A is the correct Answer

upvoted 1 times

sunil299 3 years, 10 months ago

Answer A seems correct as Replicated copies of non-searchable data are smaller than copies of searchable data, because they include only the data and not the associated index files.

upvoted 2 times

Question #2 *Topic 1*

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.

Suggested Answer: *B*
Reference:
https://docs.splunk.com/Documentation/Splunk/7.3.2/DistSearch/SHCarchitecture

Community vote distribution
A (100%)

sunil299 Highly Voted 3 years, 10 months ago
Answer is A: Replicated copies of non-searchable data are smaller than copies of searchable data, because they include only the data and not the associated index files. Increase search factor for searchable data
upvoted 7 times

bobixaka Most Recent 9 months, 4 weeks ago
Selected Answer: A
https://docs.splunk.com/Documentation/Splunk/latest/Indexer/Thesearchfactor#Search_factor
upvoted 1 times

FreshLearn 10 months ago
another question with too little context, as A depends on B and A and B depend on C.
If your search factor is already at replication factor (e.g. sf=3, rf=3) you would need to increase both and if your sf equals the current count of your search heads, you would need to increase their count as well.
Assuming that's not the case the answer can only be A) as that's required in all cases.
upvoted 1 times

bigc00p 1 year, 4 months ago
Selected Answer: A
Answer is A, why does it say B? That will only increase non searchable copies
upvoted 3 times

gsplunker 1 year, 9 months ago
Selected Answer: A
A is correct
upvoted 1 times

SplunkStreamer 2 years, 3 months ago
Selected Answer: A
Answer: A
upvoted 1 times

huu_nguyen 2 years, 7 months ago
Selected Answer: A
A is the one
upvoted 2 times

sutcocuk 2 years, 8 months ago
Selected Answer: A
A is correct
upvoted 1 times

Crash_Override 2 years, 9 months ago
Selected Answer: A

Only increasing the search factor with effect searchable data. increasing the replication factor will not effect the searchable data because it replicates non-searchable data.

upvoted 3 times

manu78 3 years, 8 months ago

A is the correct Answer

upvoted 3 times

sadhka 4 years, 2 months ago

I believe the answer should be A

upvoted 3 times

Question #3 *Topic 1*

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.

Suggested Answer: *C*

Community vote distribution

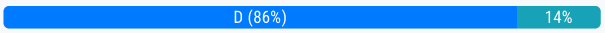

sadhka Highly Voted 4 years, 2 months ago

Yes D is the correct Answer

upvoted 9 times

YBataineh Highly Voted 4 years, 2 months ago

D is the correct Answer

upvoted 6 times

bobixaka Most Recent 10 months ago

Selected Answer: C

Before adding more search peers, you should always try to optimize the scheduled searches first.

Architecting Splunk Enterprise Deployments p.145 - point 3

upvoted 2 times

SplunkStreamer 2 years, 3 months ago

Selected Answer: D

Answer: D

upvoted 2 times

tduarte14 2 years, 6 months ago

Selected Answer: D

D is the correct answer.

upvoted 2 times

sutcocuk 2 years, 8 months ago

Selected Answer: D

D is the correct answer

upvoted 2 times

Crash_Override 2 years, 9 months ago

Selected Answer: D

most search performance issues can be addressed by adding search peers (indexers)

upvoted 6 times

manu78 3 years, 8 months ago

D is the correct Answer

upvoted 4 times

Question #4 *Topic 1*

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 sourcetype. Further investigation reveals that not all web logs 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 for 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.

Suggested Answer: *C*

Community vote distribution

Untaked 10 months ago
It's the B since they mention that the reason of the issue is that sourcetype if the only affecting the data which means that some inputs could have a wrong sourcetype name in the inputs.conf
upvoted 1 times

bobixaka 9 months, 4 weeks ago
Nope. We are talking about the same sourcetype, different parsing/format here.
upvoted 1 times

sutcocuk 2 years, 8 months ago
Selected Answer: C
C is correct
upvoted 2 times

sutcocuk 2 years, 8 months ago
C is correct
upvoted 1 times

manu78 3 years, 8 months ago
C is the correct Answer
upvoted 3 times

sadhka 4 years, 2 months ago
I think answer is B, Why the configuration of indexer and Heavy forwarder should be same.
upvoted 3 times

mker 4 years, 1 month ago
The correct answer is C.
Alternative B cannot be since the UFs cannot be configured in the props.conf and neither does it contemplate the indexers.
upvoted 6 times

RichLV 3 years, 7 months ago
Question does not specify whether other forwarders are UFs. It only mentions heavy forwarders. Could be B.
upvoted 3 times

mker 3 years, 4 months ago
For there to be a correct parsing of the data in the indexers and heavy forwarders, the same configuration must be used.
upvoted 1 times

SPLTony 1 year, 1 month ago
That's not true. Props.conf can indeed be in Universal Forwarders. For example, EVENT_BREAKER properties are ONLY applicable in props.conf on UFs.
upvoted 1 times

Question #5 *Topic 1*

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 search is locked out?

A. 300GB. After this limit, search is locked out.

B. 500GB. After this limit, search is locked out.

C. 800GB. After this limit, search is locked out.

D. Search is not locked out. Violations are still recorded.

Suggested Answer: *D*

Reference:

https://docs.splunk.com/Documentation/Splunk/7.3.2/Admin/TypesofSplunklicenses

Community vote distribution

D (100%)

SpTester Highly Voted 3 years, 5 months ago

Answer is D. Based on the document:

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

Splunk Enterprise license An Enterprise license stack with a license volume of 100 GB of data per day or more does not currently violate.

upvoted 8 times

george2t2d Most Recent 1 year, 2 months ago

D for sure

upvoted 1 times

gsplunker 1 year, 8 months ago

Selected Answer: D

Answer is D

upvoted 1 times

huu_nguyen 2 years, 7 months ago

Selected Answer: D

D is correct

upvoted 2 times

sutcocuk 2 years, 8 months ago

Selected Answer: D

D is correct

upvoted 1 times

delara 2 years, 8 months ago

D is the correct

upvoted 2 times

Question #6 *Topic 1*

What does the deployer do in a Search Head Cluster (SHC)? (Select all that apply.)

A. Distributes apps to SHC members.

B. Bootstraps a clean Splunk install for a SHC.

C. Distributes non-search related and manual configuration file changes.

D. Distributes runtime knowledge object changes made by users across the SHC.

Suggested Answer: *A*
Reference:
https://docs.splunk.com/Documentation/Splunk/7.3.2/DistSearch/SHCdeploymentoverview

Community vote distribution
A (100%)

manu78 Highly Voted 3 years, 8 months ago
A and C
upvoted 12 times

sadhka Highly Voted 4 years, 2 months ago
A and C
upvoted 7 times

Ashton_98 3 years, 11 months ago
You're correct - https://docs.splunk.com/Documentation/Splunk/7.3.2/DistSearch/PropagateSHCconfigurationchanges
upvoted 1 times

Floyda 2 years, 3 months ago
"You use the deployer to deploy configuration updates only. You cannot use it for initial configuration of the search head cluster or for version upg to the Splunk Enterprise instances that the members run on."
From the same link you shared. Only A is correct
upvoted 1 times

minombrerodrigo 1 year, 10 months ago
"Types of updates that the deployer handles
These are the specific types of updates that require the deployer:
All non-search-related updates, even those that can be configured through the CLI or Splunk Web, such as updates to indexes.conf or inputs.co

https://docs.splunk.com/Documentation/Splunk/7.3.2/DistSearch/PropagateSHCconfigurationchanges#Types_of_updates_that_the_deployer_ha
upvoted 2 times

Caro27 Most Recent 2 months, 3 weeks ago
"The deployer has these main roles:

It handles migration of app and user configurations into the search head cluster from non-cluster instances and search head pools.
It deploys baseline app configurations to search head cluster members.
It provides the means to distribute non-replicated, non-runtime configuration updates to all search head cluster members".
upvoted 1 times

bobixaka 9 months, 4 weeks ago
Selected Answer: A
Only A is correct.
B is just to confuse you by the word "bootstrap"
C is incorrect, because "You do not use the deployer to distribute search-related runtime configuration changes."
upvoted 1 times

CactiAZ 1 month ago
You have answer C written incorrectly. Answer C is written as "Distributes non-search related and manual configuration file changes" and that is exactly what the deployer is built for. The answer is A and C.

upvoted 1 times

denominator 1 year, 11 months ago

A and C, you can make manual config changes to your Search ahead cluster, like make your web.conf file to allow https for web ui

upvoted 1 times

sunil299 3 years, 10 months ago

Only A seems correct.

It deploys baseline app configurations to search head cluster members.

It provides the means to distribute non-replicated, non-runtime configuration updates to all search head cluster members. Does not copy any manual changes to other SHC

upvoted 3 times

qtygbapjpesdayazko 1 year, 3 months ago

This is the way

upvoted 1 times

demarko 3 years, 11 months ago

https://docs.splunk.com/Documentation/Splunk/8.1.0/DistSearch/PropagateSHCconfigurationchanges

upvoted 1 times

Ashton_98 4 years ago

A and D.

upvoted 2 times

Question #7 *Topic 1*

When using the props.conf LINE_BREAKER attribute to delimit multi-line events, the SHOULD_LINEMERGE attribute should be set to what?

A. Auto

B. None

C. True

D. False

Suggested Answer: *C*
Reference:
https://answers.splunk.com/answers/6926/how-to-keep-data-together-as-one-event.html

Community vote distribution
D (94%) 6%

Crash_Override Highly Voted 2 years, 9 months ago
Selected Answer: D
Should be D False
upvoted 11 times

gatundu_ Most Recent 2 months, 1 week ago
Answer is D.
upvoted 1 times

Caro27 2 months, 3 weeks ago
Totally False, check props.conf for confirmatiom
upvoted 1 times

bobixaka 9 months, 4 weeks ago
Selected Answer: D
D
upvoted 1 times

Huli7 1 year, 4 months ago
D
upvoted 1 times

emlch 1 year, 4 months ago
Selected Answer: C
If D how the mentioned multi-line event will be combined together? The answer is C. We're talking about multi-line events not single-line events. You can't work only with LINE_BREAKER
upvoted 1 times

RedYeti 1 year, 2 months ago
from props.conf documentation:

You get a significant boost to processing speed when you use LINE_BREAKER to delimit multi-line events (as opposed to using SHOULD_LINEMERGE to reassemble individual lines into multi-line events).
When using LINE_BREAKER to delimit events, SHOULD_LINEMERGE should be set to false, to ensure no further combination of delimited events occurs.
upvoted 3 times

MaryKey 1 year, 4 months ago
Selected Answer: D
D. False
upvoted 1 times

stwong 2 years ago
Seems it's A as talking about multi-line events as mentioned in props.conf doc:

SHOULD_LINEMERGE = <boolean>
* Whether or not to combine several lines of data into a single
multiline event, based on the configuration settings listed in
this subsection.
* When you set this to "true", Splunk software combines several lines of data
into a single multi-line event, based on values you configure
in the following settings.
* When you set this to "false", Splunk software does not combine lines of
data into multiline events.
* Default: true
upvoted 1 times

huu_nguyen 2 years, 7 months ago

Selected Answer: D

D is the one, guys
upvoted 2 times

manu78 3 years, 8 months ago

False is the answer
upvoted 3 times

marvinortega 3 years, 11 months ago

https://docs.splunk.com/Documentation/Splunk/8.1.0/Data/Configureeventlinebreaking#Break_the_data_stream_directly_into_real_events_with_the_LINE
upvoted 1 times

marvinortega 3 years, 11 months ago

NOTE: You get a significant boost to processing speed when you use
LINE_BREAKER to delimit multi-line events (as opposed to using
SHOULD_LINEMERGE to reassemble individual lines into multi-line events).
* When using LINE_BREAKER to delimit events, SHOULD_LINEMERGE should be set
to false, to ensure no further combination of delimited events occurs.

False is the answer.
upvoted 4 times

Question #8 *Topic 1*

Which of the following should be included in a deployment plan?

A. Business continuity and disaster recovery plans.

B. Current logging details and data source inventory.

C. Current and future topology diagrams of the IT environment.

D. A comprehensive list of stakeholders, either direct or indirect.

Suggested Answer: *D*

Reference:

https://docs.splunk.com/Documentation/CoE/ssf/Handbook/StakeholderReg

Community vote distribution

B (100%)

b5white 1 year, 3 months ago

B. It would be C except for the word 'future'.

upvoted 2 times

RedYeti 2 years, 7 months ago

Selected Answer: B

Answer B

upvoted 3 times

huu_nguyen 2 years, 7 months ago

Selected Answer: B

It's B. Architecting Splunk 8.0.1 Enterprise Deployments.pdf, page 15

upvoted 4 times

Crash_Override 2 years, 9 months ago

Selected Answer: B

Correct answer is B

upvoted 4 times

manu78 3 years, 7 months ago

B is correct

upvoted 3 times

sadhka 4 years, 2 months ago

Answer - B

upvoted 4 times

Question #9 *Topic 1*

A multi-site indexer cluster can be configured using which of the following? (Select all that apply.)

A. Via Splunk Web.

B. Directly edit SPLUNK_HOME/etc/system/local/server.conf

C. Run a splunk edit cluster-config command from the CLI.

D. Directly edit SPLUNK_HOME/etc/system/default/server.conf

Suggested Answer: *AB*
Reference:
https://docs.splunk.com/Documentation/Splunk/7.3.2/Indexer/Enableclustersindetail

Community vote distribution
BC (100%)

[Removed] Highly Voted 3 years, 8 months ago
B & C is the correct answer.
upvoted 13 times

sadhka Highly Voted 4 years, 2 months ago
B and C
Configure multisite nodes
To deploy and configure multisite cluster nodes, you must directly edit server.conf or use the CLI. You cannot use Splunk Web.
upvoted 7 times

xobeji1808 Most Recent 4 months, 4 weeks ago
A. Via Splunk Web.
C. Run a `splunk edit cluster-config` command from the CLI.

Explanation:
- **A. Via Splunk Web**: Splunk Web provides a graphical interface for configuring a multi-site indexer cluster.
- **C. Run a `splunk edit cluster-config` command from the CLI**: This CLI command allows direct configuration of the indexer cluster setup, including multi-site configurations.

Directly editing configuration files (`B` and `D`) is generally not recommended as it can lead to configuration errors or inconsistencies unless done with caution and proper understanding of Splunk's configuration management practices.
SOURCE : https://bitly.cx/8fAY7
upvoted 1 times

bobixaka 9 months, 4 weeks ago
Selected Answer: BC
https://docs.splunk.com/Documentation/Splunk/latest/Indexer/Multisitedeploymentoverview#Configure_multisite_nodes:~:text=To%20deploy%20and%2(
upvoted 2 times

gsplunker 1 year, 8 months ago
Selected Answer: BC
B&C are correct
upvoted 2 times

Crash_Override 2 years, 9 months ago
Selected Answer: BC
Correct answers are B&C
upvoted 6 times

IDM 3 years, 10 months ago
B & C is correct
Configure multisite nodes
To deploy and configure multisite cluster nodes, you must directly edit server.conf or use the CLI. You cannot use Splunk Web.

Multisite-specific configuration settings

When you deploy a multisite cluster, you configure the same settings as for single-site, along with some additional settings to specify the set of sites and the location of replicated and searchable copies across the sites.

https://docs.splunk.com/Documentation/Splunk/7.3.2/Indexer/Multisitedeploymentoverview

upvoted 3 times

[Removed] 4 years ago

https://docs.splunk.com/Documentation/Splunk/7.3.2/Indexer/Multisitedeploymentoverview

Configure with server.conf

To configure a multisite master node with server.conf, see "Configure multisite indexer clusters with server.conf".

Configure with the CLI

To configure a multisite master node with the CLI, see "Configure multisite indexer clusters with the CLI"

upvoted 2 times

Question #10 *Topic 1*

Which index-time props.conf attributes impact indexing performance? (Select all that apply.)

A. REPORT

B. LINE_BREAKER

C. ANNOTATE_PUNCT

D. SHOULD_LINEMERGE

Suggested Answer: *BD*
Reference:
https://docs.splunk.com/Documentation/Splunk/7.3.2/Data/Configureeventlinebreaking

Community vote distribution
BC (50%) BD (28%) CD (22%)

ChantreyC Highly Voted 3 years, 10 months ago
BCD - pg141-143 architecting splunk pdf
upvoted 21 times

dpharker Highly Voted 2 years, 6 months ago
Selected Answer: BC
Answers are BCD - pq 141 Architecting Splunk pdf
upvoted 9 times

marinatedcohort 1 year ago
verified in PDF, pg 141 as dpharker stated
upvoted 1 times

bobixaka Most Recent 9 months, 4 weeks ago
Selected Answer: BD
This is a very tricky question.
Answer C is questionable.

Architecting PDF pages141 and 143 states that Indexing time improves significantly by including the ANNOTATE_PUNCT parameter.

Troubleshooting PDF page 52 shows the "Great 8" rules per sourcetype will maximize the indexing performance, but they don't include the ANNOTATE_PUNCT parameter.
upvoted 1 times

Untaked 10 months, 1 week ago
I will just SAY BCD are the correct ones and left this
Annotation Processor configured

ANNOTATE_PUNCT = <boolean>
* Determines whether to index a special token starting with "punct::"
* The "punct::" key contains punctuation in the text of the event.
It can be useful for finding similar events
* If it is not useful for your dataset, or if it ends up taking
too much space in your index it is safe to disable it
* Default: true
upvoted 1 times

frappe 2 years, 4 months ago
Selected Answer: BD
Nothing in Splunk's docs specifically say that ANNOTATE_PUNCT will improve performance (it obviously will but so will a ton of other settings, and it's negligible), whereas it's consistently called out that LINE_BREAKER and SHOULD_LINEMERGE go hand in hand and will affect performance greatly.

Is the exam tricking us?

from props.conf:
* NOTE: You get a significant boost to processing speed when you use
LINE_BREAKER to delimit multi-line events (as opposed to using
SHOULD_LINEMERGE to reassemble individual lines into multi-line events).
* When using LINE_BREAKER to delimit events, SHOULD_LINEMERGE should be set
to false, to ensure no further combination of delimited events occurs.
upvoted 4 times

RedYeti 2 years, 7 months ago

Selected Answer: CD

Answers are B, C and D:
ANNOTATE_PUNCT (AP) and SHOULD_LINEMERGE (LM) which goes hand-in-hand with LINE_BREAKER (LB).
See chapter "Tune props.conf" of Architecting Splunk Enterprise Deployment. The best indexing pipelines test results are when AP and LM (so LB too) are configured.
upvoted 4 times

manu78 3 years, 7 months ago

bcd are correct
upvoted 3 times

sunil299 3 years, 10 months ago

Answer should be C and D
ANNOTATE_PUNCT = <boolean> * If it is not useful for your dataset, or if it ends up taking
too much space in your index it is safe to disable it * Default: true
upvoted 1 times

New_user 3 years, 8 months ago

Answer CD was right. 1) The REPORT option is used to order stanzas when extracting fields 2) ANNOTATE_PUNKT extracts punctuation characters from events (and doesn't influence common performance) 3) LINE_BREAKER helps to separate multi-line events to different lines (improves performance) 4) SHOULD_LINEMERGE combines lines of data to multiline events (decreases performance). Source: https://docs.splunk.com/Documentation/Splunk/8.1.2/Admin/Propsconf
upvoted 3 times

SasnycoN 2 years, 9 months ago

From what you just said we can clearly see that B is also Correct as it affects performance.
upvoted 1 times

Question #11 *Topic 1*

Which of the following are client filters available in serverclass.conf? (Select all that apply.)

A. DNS name.

B. IP address.

C. Splunk server role.

D. Platform (machine type).

Suggested Answer: *AB*

Reference:

https://docs.splunk.com/Documentation/Splunk/7.3.1/Updating/Filterclients#Define_filters_through_serverclass.conf

Community vote distribution

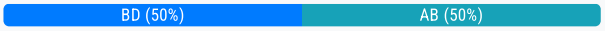

water_tank Highly Voted 3 years, 11 months ago

should be A,B and D

upvoted 15 times

manu78 Highly Voted 3 years, 7 months ago

ABD are correct

upvoted 6 times

MartinCaplan Most Recent 3 months, 1 week ago

Selected Answer: AB

ABD is correct

upvoted 1 times

bobixaka 9 months, 4 weeks ago

Selected Answer: BD

A B D are correct.

Ref: https://docs.splunk.com/Documentation/Splunk/latest/Updating/Filterclients#Types_of_filters

upvoted 2 times

Untaked 10 months, 1 week ago

Correct ones are ABD

The

filters can be based on DNS name, IP address, build number of client

machines, platform, and the clientName. If a target machine

matches the filter, then the deployment server deploys the apps and configuration

content that make up the server class to that machine.

upvoted 1 times

adamsca 1 year, 2 months ago

Selected Answer: AB

A B D is Correct

upvoted 1 times

deepali_2710 1 year, 7 months ago

ABD

Types of filters

There are three types of client filters:

Include (whitelist). Specifies clients to include, based on IP address, host name, DNS name, or client name.

Exclude (blacklist). Specifies clients to exclude, based on IP address, host name, DNS name, or client name.

Machine type. Specifies clients to include, based on machine type, such as linux-i686, linux-x86_64, and so on.

upvoted 1 times

gsplunker 1 year, 8 months ago

Selected Answer: BD

A,B & D

upvoted 1 times

giubal 1 year, 10 months ago

should be ABD

"#Server classes are essentially categories. They use filters to control
what clients they apply to, contain a set of applications, and might define
deployment server behavior for the management of those applications. The
filters can be based on DNS name, IP address, build number of client
machines, platform, and the clientName"

https://docs.splunk.com/Documentation/Splunk/8.2.5/Admin/Serverclassconf

upvoted 2 times

stwong 2 years ago

A,B,D

Types of filters

There are three types of client filters:

Include (whitelist). Specifies clients to include, based on IP address, host name, DNS name, or client name.
Exclude (blacklist). Specifies clients to exclude, based on IP address, host name, DNS name, or client name.
Machine type. Specifies clients to include, based on machine type, such as linux-i686, linux-x86_64, and so on.

https://docs.splunk.com/Documentation/Splunk/9.0.2/Updating/Filterclients

upvoted 2 times

sovip52250 2 years, 1 month ago

Selected Answer: AB

ip address
dns

upvoted 1 times

RedYeti 2 years, 7 months ago

Answers A, B and D

https://docs.splunk.com/Documentation/Splunk/8.1.3/Updating/Filterclients

upvoted 1 times

khart 3 years, 7 months ago

https://docs.splunk.com/Documentation/Splunk/8.1.3/Updating/Filterclients

upvoted 3 times

caesim 3 years, 9 months ago

correct answer should be A and B
there is only host name, ip address and dns possible to set

upvoted 1 times

Question #12 *Topic 1*

What log file would you search to verify if you suspect there is a problem interpreting a regular expression in a monitor stanza?

A. btool.log

B. metrics.log

C. splunkd.log

D. tailing_processor.log

Suggested Answer: *C*
Reference:
https://answers.splunk.com/answers/479312/how-to-edit-inputsconf-to-monitor-multiple-files-w-1.html

Community vote distribution
C (100%)

Hamiltonian Highly Voted 3 years, 4 months ago
splunkd.log
"The primary log for the Splunk server. The log is often requested by Splunk Support for troubleshooting purposes."
https://docs.splunk.com/Documentation/Splunk/8.2.1/Troubleshooting/WhatSplunklogsaboutitself
Also, metrics.log does not provide error messages or diagnostics. Troubleshooting pdf pg. 50
upvoted 7 times

sovip52250 Most Recent 2 years, 1 month ago
Selected Answer: C
splunkd.log
upvoted 2 times

Question #13 *Topic 1*

Which Splunk tool offers a health check for administrators to evaluate the health of their Splunk deployment?

A. btool

B. DiagGen

C. SPL Clinic

D. Monitoring Console

Suggested Answer: *D*

Reference:

https://docs.splunk.com/Documentation/Splunk/7.3.1/DMC/DMCoverview

Community vote distribution

D (100%)

Bianchi Highly Voted 2 years, 10 months ago

D is correct

upvoted 9 times

adamsca Most Recent 1 year, 2 months ago

Selected Answer: D

is Correct

upvoted 1 times

Dori77777 2 years ago

Selected Answer: D

is correct

upvoted 1 times

Question #14 *Topic 1*

In a four site indexer cluster, which configuration stores two searchable copies at the origin site, one searchable copy at site2, and a total of four searchable copies?

A. site_search_factor = origin:2, site1:2, total:4

B. site_search_factor = origin:2, site2:1, total:4

C. site_replication_factor = origin:2, site1:2, total:4

D. site_replication_factor = origin:2, site2:1, total:4

Suggested Answer: *D*

Reference:

https://docs.splunk.com/Documentation/Splunk/7.3.2/Indexer/Sitereplicationfactor

Community vote distribution

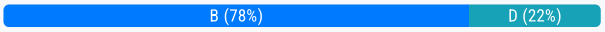

sadhka Highly Voted 4 years, 2 months ago

B should be answer

upvoted 16 times

[Removed] Highly Voted 3 years, 8 months ago

B, as we are talking about searchable copies and not replication policy

upvoted 8 times

Vidomina Most Recent 7 months, 2 weeks ago

Selected Answer: B

It's B.

D is just raw copies of data among sites, B is raw copies including tsidx files = searchable copies.

upvoted 1 times

bobixaka 9 months, 4 weeks ago

Selected Answer: B

Ref: https://docs.splunk.com/Documentation/Splunk/latest/Indexer/Sitesearchfactor#Syntax

upvoted 1 times

ginfaced 1 year, 5 months ago

It's B... This attribute specifies the per-site searchable copy policy. It is specified globally and applies to all buckets in all indexes.

https://docs.splunk.com/Documentation/Splunk/9.0.5/Indexer/Sitesearchfactor

upvoted 1 times

gsplunker 1 year, 8 months ago

Selected Answer: B

Search copy config, B is correct

upvoted 1 times

Dori77777 2 years ago

Selected Answer: D

site_search_factor = origin:2, site2:1, total:4

upvoted 2 times

Dori77777 2 years ago

B: site_search_factor = origin:2, site2:1, total:4

upvoted 3 times

frappe 2 years, 4 months ago

Selected Answer: B

The question is about searchable copies, not replication policy

upvoted 4 times

SasnycoN 2 years, 10 months ago

B is correct: https://docs.splunk.com/Documentation/Splunk/8.2.4/Indexer/Sitesearchfactor

upvoted 4 times

manu78 3 years, 7 months ago

N is correct

upvoted 1 times

sunil299 3 years, 10 months ago

Answer B: For example, if you have a three-site cluster with "site_replication_factor = origin:2, site1:1, site2:2, total:5", then, in site_search_factor, the origin value cannot exceed 2, the site1 value cannot exceed 1, the site2 value cannot exceed 2, and the total value cannot exceed 5.

upvoted 4 times

Question #15 *Topic 1*

Which Splunk Enterprise offering has its own license?

A. Splunk Cloud Forwarder

B. Splunk Heavy Forwarder

C. Splunk Universal Forwarder

D. Splunk Forwarder Management

Suggested Answer: *C*
Reference:
https://docs.splunk.com/Splexicon:Forwardinglicense

Community vote distribution
C (83%) B (17%)

huu_nguyen Highly Voted 2 years, 7 months ago

Selected Answer: C

C is the final one.

upvoted 6 times

sovip52250 2 years, 1 month ago

https://docs.splunk.com/Splexicon:Forwardinglicense#:~:text=Splunk%20Enterprise%20offers%20several%20forwarder,package%20includes%20its%20

upvoted 3 times

CactiAZ Most Recent 1 month ago

Selected Answer: C

Answer is definitely C, it's the only one that comes with its own license.

upvoted 1 times

stylox 6 months ago

B should be the answer, reason being Universal forwarder is not an enterprise.

upvoted 2 times

BRYE33 3 months, 3 weeks ago

Answer is C:

Directly from the Splunk site:
License for a forwarder, which is a Splunk Enterprise instance that forwards data to another Splunk Enterprise server or to a third-party system.

Splunk Enterprise offers several forwarder licensing options:

The universal forwarder package includes its own license. The license is enabled or applied automatically. This license allows forwarding but not indexing of unlimited data, and also enables security on the forwarder so that users must supply a user name and password to access it. The heavy forwarder should have access to an Enterprise license stack if you plan to perform indexing on the forwarder or to enable authentication on the forwarder.

upvoted 1 times

deepali_2710 1 year, 7 months ago

C. Splunk Universal Forwarder

upvoted 2 times

Hemnaath 1 year, 9 months ago

Correct answer is "C"
Under Data Administration course --> Module3 Forwarder Configuration --> Understanding Universal forwarder --> It is mentioned that UF provided a separate installation binary with built in license (no limit)

upvoted 1 times

sovip52250 2 years, 1 month ago

Selected Answer: B

Splunk Heavy Forwarder

upvoted 2 times

olibee 2 years, 2 months ago

B is the correct answer. Only HF and the Enterprise Instance need a license. No "normal" forwarder would need it.

upvoted 3 times

rahendri 2 years, 2 months ago

the answer is B as there is no license requirement for an UF, but a HF needs a license, which can be set in the license menu.

upvoted 4 times

RedYeti 2 years, 7 months ago

Selected Answer: C

Answer C

upvoted 4 times

Splunky_Manny 3 years, 5 months ago

I agree. C is the correct answer. You don't a license to install the UF

upvoted 2 times

[Removed] 3 years, 8 months ago

C is correct

upvoted 3 times

Question #16 *Topic 1*

Which component in the splunkd.log will log information related to bad event breaking?

A. Audittrail

B. EventBreaking

C. IndexingPipeline

D. AggregatorMiningProcessor

Suggested Answer: *D*

Reference:

https://answers.splunk.com/answers/141721/error-in-splunkd-log-breaking-event-because-limit-of-256-has-been-exceeded.html

Community vote distribution

D (100%)

delara Highly Voted 2 years, 8 months ago

D, page 72 troubleshooting guide

upvoted 6 times

Bianchi Highly Voted 2 years, 10 months ago

D is correct, as per : https://docs.splunk.com/Documentation/Splunk/8.2.4/Data/Resolvedataqualityissues

upvoted 5 times

marinatedcohort 1 year ago

perfect reference

upvoted 1 times

qtygbapjpesdayazko Most Recent 1 year, 3 months ago

D. AggregatorMiningProcessor

Splunk Name, spelunking is the hobby of exploring caves and mines. Splunking, then, is the exploration of information caves and the mining of data.

upvoted 2 times

gorasz 1 year, 11 months ago

Selected Answer: D

D is the answer

upvoted 2 times

SasnycoN 2 years, 10 months ago

D is correct

upvoted 2 times

Question #17 *Topic 1*

Which Splunk server role regulates the functioning of indexer cluster?

A. Indexer

B. Deployer

C. Master Node

D. Monitoring Console

Suggested Answer: *C*
Reference:
https://docs.splunk.com/Documentation/Splunk/7.3.1/Deploy/Indexercluster

Community vote distribution
C (100%)

SasnycoN Highly Voted 2 years, 10 months ago

Selected Answer: C

Answer C - Master (now manager) node: https://docs.splunk.com/Documentation/Splunk/8.2.4/Deploy/Indexercluster
https://docs.splunk.com/Documentation/Splunk/8.2.4/Indexer/Enablethemanagernode
upvoted 8 times

SasnycoN Highly Voted 2 years, 10 months ago

Answer c - Master (now manager) node: https://docs.splunk.com/Documentation/Splunk/8.2.4/Deploy/Indexercluster
https://docs.splunk.com/Documentation/Splunk/8.2.4/Indexer/Enablethemanagernode
upvoted 5 times

Question #18 *Topic 1*

When adding or rejoining a member to a search head cluster, the following error is displayed:
Error pulling configurations from the search head cluster captain; consider performing a destructive configuration resync on this search head cluster member.
What corrective action should be taken?

A. Restart the search head.

B. Run the splunk apply shcluster-bundle command from the deployer.

C. Run the clean raft command on all members of the search head cluster.

D. Run the splunk resync shcluster-replicated-config command on this member.

Suggested Answer: *B*

Community vote distribution

D (100%)

sadhka Highly Voted 4 years, 2 months ago

D
https://community.splunk.com/t5/Deployment-Architecture/How-to-resolve-error-quot-Error-pulling-configurations-from-the/m-p/354231
upvoted 16 times

deepali_2710 Most Recent 1 year, 7 months ago

Error pulling configurations from the search head cluster captain; consider performing a destructive configuration resync on this search head cluster member.
The message also appears in the member's splunkd.logfile.

If this message appears, it means that the member is unable to update its configuration through the configuration change delta and must apply the entire configuration tarball. It does not do this automatically. Instead, it waits for your intervention.

You must then initiate the process of downloading and applying the tarball by running this CLI command on the member:

splunk resync shcluster-replicated-config
upvoted 3 times

qtygbapjpesdayazko 1 year, 3 months ago

this is the way
upvoted 2 times

lzng3r 1 year, 7 months ago

Selected Answer: D
https://docs.splunk.com/Documentation/Splunk/latest/DistSearch/HowconfrepoworksinSHC#Perform_a_manual_resync
upvoted 2 times

ozii 2 years, 8 months ago

D is correct
upvoted 3 times

delara 2 years, 8 months ago

D, page 196 Cluster admin
upvoted 2 times

delara 2 years, 9 months ago

D is the correct
upvoted 1 times

Crash_Override 2 years, 9 months ago

Selected Answer: D
D is the correct answer
upvoted 3 times

Hamiltonian 3 years, 4 months ago

Answer is D.

See Clustering slides 196 and 205. Also:

https://docs.splunk.com/Documentation/Splunk/latest/DistSearch/HowconfrepoworksinSHC

upvoted 2 times

Question #19 *Topic 1*

Which of the following commands is used to clear the KV store?

A. splunk clean kvstore

B. splunk clear kvstore

C. splunk delete kvstore

D. splunk reinitialize kvstore

Suggested Answer: *A*
Reference:
https://answers.splunk.com/answers/237859/can-i-delete-all-data-from-a-kv-store-at-once.html

Community vote distribution
A (100%)

SasnycoN Highly Voted 2 years, 10 months ago
Answer is A: https://docs.splunk.com/Documentation/Splunk/8.2.4/Admin/ResyncKVstore
upvoted 6 times

qtygbapjpesdayazko Most Recent 1 year, 3 months ago
Selected Answer: A
the command is
splunk clean kvstore
upvoted 1 times

asashima 2 years, 12 months ago
Answer is A.
See Clustering slides 222.
upvoted 4 times

Question #20 *Topic 1*

Indexing is slow and real-time search results are delayed in a Splunk environment with two indexers and one search head. There is ample CPU and memory available on the indexers. Which of the following is most likely to improve indexing performance?

A. Increase the maximum number of hot buckets in indexes.conf

B. Increase the number of parallel ingestion pipelines in server.conf

C. Decrease the maximum size of the search pipelines in limits.conf

D. Decrease the maximum concurrent scheduled searches in limits.conf

Suggested Answer: *D*

Community vote distribution

B (100%)

M_K_S Highly Voted 4 years ago

Should be B

upvoted 9 times

Crash_Override Highly Voted 2 years, 9 months ago

Selected Answer: B

B is correct answer

upvoted 7 times

ginfaced Most Recent 1 year, 5 months ago

I went with D as the parallelization features are intended for customers with excess CPU cores and I/O capacity to leverage their hardware for improved performance across the indexing tier.

upvoted 1 times

Splunky_Manny 3 years, 5 months ago

Parallel ingestion increases performance https://conf.splunk.com/files/2016/slides/harnessing-performance-and-scalability-with-parallelization.pdf

upvoted 3 times

manu78 3 years, 7 months ago

B is correct

upvoted 6 times

Question #21 *Topic 1*

The guidance Splunk gives for estimating size on for syslog data is 50% of original data size. How does this divide between files in the index?

A. rawdata is: 10%, tsidx is: 40%

B. rawdata is: 15%, tsidx is: 35%

C. rawdata is: 35%, tsidx is: 15%

D. rawdata is: 40%, tsidx is: 10%

Suggested Answer: *B*

Reference:

https://answers.splunk.com/answers/147951/what-is-the-compression-ratio-of-raw-data-in-splunk.html

Community vote distribution

B (100%)

SasnycoN Highly Voted 2 years, 10 months ago

Selected Answer: B

Answer is B: https://docs.splunk.com/Documentation/Splunk/8.2.4/Capacity/Estimateyourstoragerequirements

upvoted 6 times

asashima Most Recent 2 years, 12 months ago

Answer is B.

See Clustering slides 27.

upvoted 2 times

Question #22 *Topic 1*

A three-node search head cluster is skipping a large number of searches across time. What should be done to increase scheduled search capacity on the search head cluster?

A. Create a job server on the cluster.

B. Add another search head to the cluster.

C. server.conf captain_is_adhoc_searchhead = true.

D. Change limits.conf value for max_searches_per_cpu to a higher value.

Suggested Answer: *D*

Community vote distribution

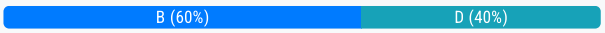

rodrigok Highly Voted 2 years, 5 months ago

correct is B: cluster-admin -145

upvoted 12 times

bobixaka Most Recent 9 months, 4 weeks ago

Selected Answer: D

Could be Answer D.

Ref: https://lantern.splunk.com/Splunk_Platform/Product_Tips/Searching_and_Reporting/Reducing_skipped_searches#:~:text=The%20total%20maximum

upvoted 2 times

bobixaka 9 months, 4 weeks ago

Could be also B

Ref:

https://docs.splunk.com/Documentation/Splunk/latest/DistSearch/SHCsystemrequirements#:~:text=You%20can%20optionally%20add%20more%20me

upvoted 1 times

willsy 2 years, 1 month ago

D

Overall search quota. This quota determines the maximum number of historical searches (combined scheduled and ad hoc) that the cluster can run concurrently. This quota is configured with max_searches_per_cpu and related settings in limits.conf

https://docs.splunk.com/Documentation/Splunk/8.2.4/DistSearch/SHCjobscheduling

upvoted 1 times

samme 2 years ago

Jacking up the max_searches_per_cpu doesn't always solve the problem. There a limit to this strategy. Once you're out of cpu, its going to have large queue and skipped searches.

My answer would be B or ensure your scheduled reports are scheduled for exact same time.

upvoted 5 times

sovip52250 2 years, 1 month ago

Selected Answer: B

A three-node search head cluster is skipping a large

upvoted 3 times

SasnycoN 2 years, 10 months ago

Should be D: https://docs.splunk.com/Documentation/Splunk/8.2.4/DistSearch/SHCjobscheduling

upvoted 3 times

sunil299 3 years, 10 months ago

Answer is C i guess.

The SHC is over-tasked with too many scheduled searches, resulting in a high skip ratio

https://aditumpartners.com/troubleshooting-splunk-search-head-clusters/

upvoted 1 times

New_user 3 years, 8 months ago

If you switch the captain_is_adhoc_searchhead to true, it will stop all the scheduled searches on the captain SH. So, D is better answer

upvoted 8 times

Bianchi 2 years, 7 months ago

This,actually , will be worse

upvoted 2 times

Question #23 *Topic 1*

The frequency in which a deployment client contacts the deployment server is controlled by what?

A. polling_interval attribute in outputs.conf

B. phoneHomeIntervalInSecs attribute in outputs.conf

C. polling_interval attribute in deploymentclient.conf

D. phoneHomeIntervalInSecs attribute in deploymentclient.conf

Suggested Answer: *D*

Reference:

https://docs.splunk.com/Documentation/SplunkCloud/7.2.7/RESTREF/RESTdeploy

Community vote distribution

SasnycoN Highly Voted 2 years, 10 months ago

Confirming answer D: https://docs.splunk.com/Documentation/Splunk/8.2.4/Admin/Deploymentclientconf

upvoted 11 times

bobixaka Most Recent 9 months, 4 weeks ago

Selected Answer: D

Ref:

https://docs.splunk.com/Documentation/Splunk/latest/Admin/Deploymentclientconf#:~:text=phoneHomeIntervalInSecs%20%3D%20%3Cdecimal%3E

upvoted 1 times

qtygbapjpesdayazko 1 year, 5 months ago

Selected Answer: D

phoneHomeIntervalInSecs = <decimal>

* How frequently, in seconds, this deployment client should

check for new content.

* Fractional seconds are allowed.

* Default: 60.

upvoted 1 times

asashima 2 years, 12 months ago

Answer is D.

upvoted 4 times

Question #24 *Topic 1*

To activate replication for an index in an indexer cluster, what attribute must be configured in indexes.conf on all peer nodes?

A. repFactor = 0

B. replicate = 0

C. repFactor = auto

D. replicate = auto

Suggested Answer: *C*
Reference:
https://docs.splunk.com/Documentation/Splunk/7.3.1/Indexer/Configurethepeerindexes

Community vote distribution
C (100%)

SasnycoN Highly Voted 2 years, 10 months ago
Confirming answer C: https://docs.splunk.com/Documentation/Splunk/8.2.4/Indexer/Configurethepeerindexes
upvoted 10 times

bobixaka Most Recent 9 months, 4 weeks ago
Selected Answer: C
Ref:
https://docs.splunk.com/Documentation/Splunk/latest/Admin/Indexesconf#:~:text=metadata%0A%20%20was%20removed.-,repFactor%20%3D%200%7Ca*%20Valid%20only%20for
upvoted 1 times

minombrerodrigo 1 year, 10 months ago
Answer is C
Page 80 Cluster-Admin PDF
upvoted 1 times

asashima 2 years, 12 months ago
Answer is C.
See Clustering slides 86.
upvoted 3 times

Question #25 *Topic 1*

Which of the following clarification steps should be taken if apps are not appearing on a deployment client? (Select all that apply.)

A. Check serverclass.conf of the deployment server.

B. Check deploymentclient.conf of the deployment client.

C. Check the content of SPLUNK_HOME/etc/apps of the deployment server.

D. Search for relevant events in splunkd.log of the deployment server.

Suggested Answer: *ABC*

Reference:

https://answers.splunk.com/answers/177021/why-is-deployment-client-not-picking-up-changes-to.html

Community vote distribution

ABD (83%) AB (17%)

sadhka Highly Voted 4 years, 2 months ago

A,B,D - There is no link between the etc/apps folder and deployment apps on the deployment server.

upvoted 23 times

New_user 3 years, 8 months ago

You're right, apps are downloaded from "deployment-apps" folder, not from "apps". But checking events in splunkd.log is also worthless. I think answer is AB

upvoted 3 times

SasnycoN 2 years, 10 months ago

D should be also correct. As per page 89 of Troubleshooting PDF we can see that there are useful error messages that you can see in splunkd.log.

upvoted 2 times

Hamiltonian 3 years, 4 months ago

I think this is the best answer. You can check serverclass.conf to verify that the client is associated with serverclass in which the app is in. You can also check deploymentclient.conf to make sure that the client is always receiving apps from the deployment server (by default should be so). Finally you can check splunkd.log for any errors or warnings with component DeploymentApplication (see Troubleshooting pdf pg 86)

upvoted 4 times

gsplunker Most Recent 1 year, 8 months ago

Selected Answer: ABD

C is not correct because apps that are to be deployed will be in deployment-apps folder in deployment server and not apps folder

upvoted 3 times

minombrerodrigo 1 year, 10 months ago

Selected Answer: ABD

C is incorrect because the folder where the apps are located in the DEPLOYMENT SERVER is SPLUNK_HOME/etc/deployment-apps

upvoted 3 times

Serderapoksiljo 1 year, 10 months ago

ABD is only logical answer

upvoted 1 times

brettw 2 years, 2 months ago

C is incorrect because apps deployed from the DS are located in SPLUNK_HOME/etc/deployment-apps

upvoted 2 times

RedYeti 2 years, 7 months ago

Selected Answer: ABD

Answers A, B and D

upvoted 2 times

bobixaka 10 months ago

D doesn't look to be correct, because the log should be reviewed on the client, not the server.
You can see if the client connects to the server successfully in the client's log, not the server's.
upvoted 1 times

RedYeti 2 years, 7 months ago

Selected Answer: ABD

Answers A, B and D as "apps" folder contain application for the Deployment Server. The applications for the clients are in "deployment-apps" folder.
upvoted 1 times

Crash_Override 2 years, 9 months ago

Selected Answer: ABD

Answer is abd
upvoted 4 times

delara 2 years, 9 months ago

Selected Answer: AB

the apps in deployment server are in deployment-app, not in app
upvoted 3 times

SasnycoN 2 years, 10 months ago

Selected Answer: ABD

Answers A, B, D. Answer D Is included as per page 89 of the Troubleshooting PDF "Example: Deployment Error Messages".
upvoted 2 times

sovip52250 2 years, 1 month ago

splunkd.log very usefull
https://community.splunk.com/t5/Deployment-Architecture/How-can-I-find-out-what-the-deployment-errors-are/m-p/325801
upvoted 1 times

bobixaka 10 months ago

Nope. D is incorrect, because the log should be reviewed on the CLIENT, not the SERVER.
This page shows messages like "*** INFO DC:DeploymentClient ***" "*** WARN DeployedApplication ***", etc...
upvoted 1 times

Question #26 *Topic 1*

Which of the following security options must be explicitly configured (i.e. which options are not enabled by default)?

A. Data encryption between Splunk Web and splunkd.

B. Certificate authentication between forwarders and indexers.

C. Certificate authentication between Splunk Web and search head.

D. Data encryption for distributed search between search heads and indexers.

Suggested Answer: *B*

Community vote distribution

scostic Highly Voted 3 years, 9 months ago

B and C https://docs.splunk.com/Documentation/Splunk/latest/Security/AboutsecuringyourSplunkconfigurationwithSSL

upvoted 12 times

jackac 2 years, 11 months ago

it actually should just be B, Splunk Web to search head is SSL secured by default but forwarder to indexer is NOT by default. The table in the link above shows this.

upvoted 3 times

SasnycoN 2 years, 10 months ago

Answer C is not for the SSL but for the Certificate authentication. Encryption between Splunk Web and SH is enabled by default but NOT the "Certificate Authentication" which is the example in C. In fact there is no single case where the Certificate authentication is enabled by default.

upvoted 4 times

Proctor Highly Voted 2 years, 1 month ago

Just FYI - in the real exam, this is not a multiple choice question. Only one answer is accepted.

upvoted 8 times

BRYE33 3 months, 2 weeks ago

So which answer would you choose @Proctor?

upvoted 1 times

CactiAZ Most Recent 1 month ago

Selected Answer: B

It's only answer B. A and C are not valid answers because they mention internal processes that occur within a single box. D is valid but is already turned on by default.

upvoted 1 times

bobixaka 9 months, 4 weeks ago

Selected Answer: C

B and C

Ref:

https://docs.splunk.com/Documentation/Splunk/latest/Security/AboutsecuringyourSplunkconfigurationwithSSL#Methods_to_secure_the_Splunk_platform

upvoted 1 times

qtygbapjpesdayazko 1 year, 5 months ago

Selected Answer: B

B and C

upvoted 1 times

RedYeti 2 years, 7 months ago

Selected Answer: C

Answers B and C.

Data encryption is enabled everywhere by default except from Forwarders to Indexers, between Indexers and from browser to Splunk Web.

In the other hand, certificate authentication is never enabled by default anywhere.

upvoted 1 times

SasnycoN 2 years, 10 months ago

Answers B and C should be valid both according to :
https://docs.splunk.com/Documentation/Splunk/latest/Security/AboutsecuringyourSplunkconfigurationwithSSL

upvoted 3 times

matsumo 3 years, 1 month ago

This question was a single answer. I think C is correct.

upvoted 2 times

RedYeti 2 years, 7 months ago

It's written "which options are not enabled by default"

upvoted 1 times

Question #27 *Topic 1*

Which of the following artifacts are included in a Splunk diag file? (Select all that apply.)

A. OS settings.

B. Internal logs.

C. Customer data.

D. Configuration files.

Suggested Answer: *BD*
Reference:
https://docs.splunk.com/Documentation/Splunk/7.3.2/Troubleshooting/Generateadiag

Community vote distribution
AB (67%) BD (33%)

sunil299 Highly Voted 3 years, 10 months ago
Answer is A,B,D
collects basic information about your Splunk platform instance, including Splunk platform configuration details. It gathers information, such as server specs, OS version, file system, and current open connections, from the machine running the Splunk platform.
upvoted 19 times

CactiAZ Most Recent 1 month ago
ABD is the answer
upvoted 1 times

srek3502 1 year, 1 month ago
https://docs.splunk.com/Documentation/Splunk/9.1.1/Troubleshooting/Generateadiag
A diag file provides a snapshot of the configurations and logs from the Splunk software along with select information about the platform instance. The diag collection process gathers information such as server specifications, operating system (OS) version, file system information, and current network connections. A diag collection also includes the contents of the $SPLUNK_HOME installation path, such as app configurations, internal log files, and index metadata.
upvoted 1 times

deepali_2710 1 year, 7 months ago
• A. OS settings: The diag file includes a snapshot of the current operating system settings, which can be useful in diagnosing issues related to hardware, networking, or system performance.

• B. Internal logs: The diag file includes logs from various internal components of Splunk, such as the indexer, search head, and deployment server. These logs can be used to diagnose issues related to indexing, searching, or configuration management.

• D. Configuration files: The diag file includes copies of various configuration files used by Splunk, such as server.conf, inputs.conf, and outputs.conf. These files can be used to verify the current configuration settings, and to troubleshoot issues related to data ingestion, forwarding, or search.
upvoted 1 times

lzng3r 1 year, 7 months ago
Selected Answer: AB
A,B,D - Troubleshooting.pdf
upvoted 1 times

KiranVM 1 year, 8 months ago
A, B & D
OS settings, internal logs, configuration files
upvoted 1 times

Vale5M 1 year, 8 months ago
A, B, D Troubleshooting.pdf slide 24
upvoted 2 times

minombrerodrigo 1 year, 10 months ago

ABD are correct.

"The diag collection process gathers information such as server specifications, operating system (OS) version, file system information, and current network connections. A diag collection also includes the contents of the $SPLUNK_HOME installation path, such as app configurations, internal log files, and index metadata."

https://docs.splunk.com/Documentation/Splunk/9.0.3/Troubleshooting/Generateadiag#About_diag

upvoted 3 times

sovip52250 2 years, 1 month ago

Selected Answer: BD

Diag contents

Primarily, a diag contains server logs, from $SPLUNK_HOME/var/log/splunk and $SPLUNK_HOME/var/log/introspection, and the configuration files, from $SPLUNK_HOME/etc.

https://docs.splunk.com/Documentation/Splunk/9.0.1/Troubleshooting/Generateadiag

upvoted 1 times

Bianchi 2 years, 7 months ago

Selected Answer: AB

ABD are correct

upvoted 1 times

huu_nguyen 2 years, 7 months ago

ABD. Troubleshooting guide, p24

upvoted 4 times

Question #28 *Topic 1*

Which command will permanently decommission a peer node operating in an indexer cluster?

A. splunk stop -f

B. splunk offline -f

C. splunk offline --enforce-counts

D. splunk decommission --enforce counts

Suggested Answer: *C*
Reference:
https://docs.splunk.com/Documentation/Splunk/7.3.2/Indexer/Takeapeeroffline

Community vote distribution
C (100%)

srek3502 1 year, 1 month ago
Answer C splunk offline --enforce-counts
Cluster Admin pdf -> pg 93
upvoted 1 times

KiranVM 1 year, 8 months ago
Selected Answer: C
splunk offline --enforce-counts
upvoted 2 times

RedYeti 2 years, 7 months ago
Selected Answer: C
Answer C

https://docs.splunk.com/Documentation/Splunk/latest/Indexer/Takeapeeroffline#The_splunk_offline_command
upvoted 2 times

asashima 2 years, 11 months ago
Answer is C
splunk offline --enforce-counts. Used to remove a peer permanently from the cluster. Also known as the "enforce-counts offline" command.
upvoted 4 times

marvinortega 3 years, 11 months ago
splunk offline --enforce-counts. Used to remove a peer permanently from the cluster. Also known as the "enforce-counts offline" command.
https://docs.splunk.com/Documentation/Splunk/8.1.0/Indexer/Takeapeeroffline#The_splunk_offline_command
upvoted 4 times

Question #29 *Topic 1*

Which CLI command converts a Splunk instance to a license slave?

A. splunk add licenses

B. splunk list licenser-slaves

C. splunk edit licenser-localslave

D. splunk list licenser-localslave

Suggested Answer: *C*
Reference:
https://docs.splunk.com/Documentation/Splunk/7.3.2/Admin/LicenserCLIcommands

Community vote distribution
C (100%)

RedYeti Highly Voted 2 years, 7 months ago
Selected Answer: C
Answer C

https://docs.splunk.com/Documentation/Splunk/latest/Admin/LicenserCLIcommands
upvoted 5 times

SPLTony Most Recent 1 year, 1 month ago
Answer C
As of today and according to the documentation, it would be
splunk edit licenser-localpeer

https://docs.splunk.com/Documentation/Splunk/latest/Admin/LicenserCLIcommands
upvoted 2 times

qtygbapjpesdayazko 1 year, 3 months ago
The correct is C, full command:
splunk edit licenser-localslave -master_uri https://YourLicenseMaster:8089
upvoted 1 times

ashutoshab 1 year, 7 months ago
Answer C
upvoted 1 times

SasnycoN 2 years, 10 months ago
Answer C Confirmed as per slide 14 of Cluster Administration PDF
upvoted 4 times

Question #30 *Topic 1*

Splunk Enterprise platform instrumentation refers to data that the Splunk Enterprise deployment logs in the _introspection index. Which of the following logs are included in this index? (Select all that apply.)

A. audit.log

B. metrics.log

C. disk_objects.log

D. resource_usage.log

Suggested Answer: *CD*
Reference:
https://docs.splunk.com/Documentation/Splunk/7.3.1/Troubleshooting/Abouttheplatforminstrumentationframework

Community vote distribution
CD (100%)

RedYeti Highly Voted 2 years, 7 months ago

Selected Answer: CD

Answers C and D

https://docs.splunk.com/Documentation/Splunk/latest/Troubleshooting/Abouttheplatforminstrumentationframework

upvoted 5 times

Bianchi Most Recent 2 years, 7 months ago

Selected Answer: CD

https://docs.splunk.com/Documentation/Splunk/8.2.5/Troubleshooting/Whatdatagetslogged

CD

upvoted 3 times

sutcocuk 2 years, 9 months ago

C&D

https://docs.splunk.com/Documentation/Splunk/8.2.5/Troubleshooting/Whatdatagetslogged

upvoted 3 times

Question #31 *Topic 1*

Which of the following can a Splunk diag contain?

A. Search history, Splunk users and their roles, running processes, indexed data

B. Server specs, current open connections, internal Splunk log files, index listings

C. KV store listings, internal Splunk log files, search peer bundles listings, indexed data

D. Splunk platform configuration details, Splunk users and their roles, current open connections, index listings

Suggested Answer: *B*

Reference:

https://splunkonbigdata.com/2018/10/01/splunk-diag/

Community vote distribution

B (100%)

qtygbapjpesdayazko 1 year, 5 months ago

Selected Answer: B

B. Server specs, current open connections, internal Splunk log files, index listings

upvoted 1 times

RedYeti 2 years, 7 months ago

Selected Answer: B

Answer B

https://docs.splunk.com/Documentation/Splunk/8.2.6/Troubleshooting/Generateadiag

in "About diag"

upvoted 4 times

huu_nguyen 2 years, 7 months ago

B is the one. Splunk diag will never collect personal data including: indexed data, user data, etc

upvoted 2 times

Question #32 *Topic 1*

Which of the following are true statements about Splunk indexer clustering?

A. All peer nodes must run exactly the same Splunk version.

B. The master node must run the same or a later Splunk version than search heads.

C. The peer nodes must run the same or a later Splunk version than the master node.

D. The search head must run the same or a later Splunk version than the peer nodes.

Suggested Answer: *B*
Reference:
https://answers.splunk.com/answers/760348/search-head-version-compatibility.html

Community vote distribution
A (100%)

khart Highly Voted 3 years, 7 months ago
slide 22 from the Cluster Admin ppt - A,B,D
upvoted 17 times

not_another_user_007 3 years, 2 months ago
In the cluster admin guide it refers the master as the manager node.
upvoted 2 times

CactiAZ Most Recent 1 month ago
A, B, and D
upvoted 1 times

lmak 10 months, 1 week ago
The true statements about Splunk indexer clustering are: A and C
a. All peer nodes must run exactly the same Splunk version. Peer nodes within a Splunk indexer cluster must be running the same Splunk version to ensure compatibility and proper communication within the cluster.
c. The peer nodes must run the same or a later Splunk version than the master node.
Peer nodes can run the same or a later Splunk version than the master node. It's important for peer nodes to be compatible with the master node's version to ensure proper functionality
upvoted 1 times

qtygbapjpesdayazko 1 year, 5 months ago
Selected Answer: A
The manager node must run the same or a later version than the peer nodes and
search heads.
The search heads must run the same or a later version than the peer node.
All peer nodes must run EXACTLY the same version.
upvoted 2 times

deepali_2710 1 year, 7 months ago
A. All peer nodes must run exactly the same Splunk version.
B. The master node must run the same or a later Splunk version than search heads.
D. The search head must run the same or a later Splunk version than the peer nodes.
upvoted 3 times

qtygbapjpesdayazko 1 year, 3 months ago
this is the way
upvoted 1 times

RedYeti 2 years, 7 months ago
Selected Answer: A
Answers A, B and D:
https://docs.splunk.com/Documentation/Splunk/latest/Indexer/Systemrequirements

A. All peer nodes must run exactly the same Splunk version.
(chapter "Compatibility between peer nodes")

B. The master node must run the same or a later Splunk version than search heads.
(chapter "Splunk Enterprise version compatibility")

Answer C is not possible as it's contradictory with answer B.

D. The search head must run the same or a later Splunk version than the peer nodes.
(chapter "Compatibility between peer nodes and search heads")
upvoted 2 times

sovip52250 2 years, 1 month ago
this question is asking about indexer clustering..not compatibility between peer node
upvoted 1 times

sutcocuk 2 years, 9 months ago
It's A,B & D
upvoted 2 times

[Removed] 3 years, 8 months ago
A,B,D : Interoperability between the various types of cluster nodes is subject to strict compatibility requirements. In brief:

The master node must run the same or a later version from the peer nodes and search heads.
The search heads must run the same or a later version from the peer nodes.
The peer nodes must all run exactly the same version, down to the maintenance level.
upvoted 3 times

IDM 3 years, 10 months ago
A and B
https://docs.splunk.com/Documentation/Splunk/8.0.6/Indexer/Systemrequirements#:~:text=All%20peer%20nodes%20must%20run,some%20peer%20nod
upvoted 2 times

marvinortega 3 years, 11 months ago
Correction, A&D
upvoted 1 times

IDM 3 years, 10 months ago
Compatibility between peer nodes and search heads
The peer nodes and search heads can run different versions from each other. The search heads must run the same or a later version from the peer nodes.
upvoted 2 times

marvinortega 3 years, 11 months ago
A,B,D
https://docs.splunk.com/Documentation/Splunk/8.0.6/Indexer/Systemrequirements#Splunk_Enterprise_version_compatibility
upvoted 2 times

IDM 3 years, 10 months ago
A, B not D
https://docs.splunk.com/Documentation/Splunk/8.0.6/Indexer/Systemrequirements#:~:text=All%20peer%20nodes%20must%20run,some%20peer%20r
upvoted 2 times

sadhka 4 years, 2 months ago
A is the correct answer.
https://docs.splunk.com/Documentation/Splunk/8.0.6/Indexer/Systemrequirements#:~:text=All%20peer%20nodes%20must%20run,some%20peer%20nod
upvoted 2 times

Question #33 *Topic 1*

A customer plans to ingest 600 GB of data per day into Splunk. They will have six concurrent users, and they also want high data availability and high search performance. The customer is concerned about cost and wants to spend the minimum amount on the hardware for Splunk. How many indexers are recommended for this deployment?

A. Two indexers not in a cluster, assuming users run many long searches.

B. Three indexers not in a cluster, assuming a long data retention period.

C. Two indexers clustered, assuming high availability is the greatest priority.

D. Two indexers clustered, assuming a high volume of saved/scheduled searches.

Suggested Answer: *D*
Reference:
https://docs.splunk.com/Documentation/Splunk/7.3.2/Capacity/Summaryofperformancerecommendations

Community vote distribution
C (100%)

sadhka Highly Voted 4 years, 2 months ago
Answer is C
upvoted 12 times

sovip52250 2 years, 1 month ago
https://docs.splunk.com/Documentation/Splunk/9.0.1/Capacity/Summaryofperformancerecommendations#:~:text=An%20indexer%20that%20meets%
upvoted 2 times

qtygbapjpesdayazko 1 year, 5 months ago
1 Search Head,
2 Indexers
upvoted 1 times

bobixaka Most Recent 10 months ago
Selected Answer: C
The answer should be C.
Scheduled searches have nothing to do with indexing and availability.
upvoted 1 times

MaryKey 1 year, 4 months ago
Selected Answer: C
The answer is C
upvoted 1 times

Redtonyeah 2 years, 8 months ago
Selected Answer: C
Answer is C
upvoted 2 times

manu78 3 years, 7 months ago
C is ok
upvoted 3 times

marvinortega 3 years, 11 months ago
Agreed with answer C.
upvoted 3 times

marvinortega 3 years, 11 months ago
"Data availability. An indexer is always available to handle incoming data, and the indexed data is available for searching."
https://docs.splunk.com/Documentation/Splunk/8.1.0/Indexer/Aboutclusters#About_indexer_clusters_and_index_replication

It means IDXC
upvoted 1 times

Question #34 *Topic 1*

To reduce the captain's work load in a search head cluster, what setting will prevent scheduled searches from running on the captain?

A. adhoc_searchhead = true (on all members)

B. adhoc_searchhead = true (on the current captain)

C. captain_is_adhoc_searchhead = true (on all members)

D. captain_is_adhoc_searchhead = true (on the current captain)

Suggested Answer: *D*

Reference:

https://docs.splunk.com/Documentation/Splunk/7.3.1/DistSearch/Adhocclustermember

Community vote distribution

C (100%)

manu78 Highly Voted 3 years, 7 months ago

C is correct

upvoted 8 times

qtygbapjpesdayazko Most Recent 1 year, 3 months ago

Selected Answer: C

To designate the captain as an ad hoc search head, set the captain_is_adhoc_searchhead attribute in server.conf on each member:

upvoted 3 times

inkedia3 1 year, 7 months ago

Selected Answer: C

Cluster Admin Page 152

upvoted 3 times

deepali_2710 1 year, 7 months ago

To designate the captain as an ad hoc search head, set the captain_is_adhoc_searchhead attribute in server.conf on each member:

[shclustering]

captain_is_adhoc_searchhead = true

upvoted 2 times

RedYeti 2 years, 7 months ago

Selected Answer: C

Answer C

https://docs.splunk.com/Documentation/Splunk/latest/DistSearch/Adhocclustermember

upvoted 1 times

Redtonyeah 2 years, 8 months ago

Selected Answer: C

Answer is C

upvoted 2 times

sunil299 3 years, 10 months ago

To designate the captain as an ad hoc search head, set the captain_is_adhoc_searchhead attribute in server.conf on each member

upvoted 1 times

[Removed] 3 years, 10 months ago

https://docs.splunk.com/Documentation/Splunk/7.3.1/DistSearch/Adhocclustermember

c

upvoted 2 times

sadhka 4 years, 2 months ago

C is the answer

upvoted 4 times

Question #35 *Topic 1*

At which default interval does metrics.log generate a periodic report regarding license utilization?

A. 10 seconds

B. 30 seconds

C. 60 seconds

D. 300 seconds

Suggested Answer: *B*
Reference:
https://docs.splunk.com/Documentation/Splunk/7.3.2/Troubleshooting/Aboutmetricslog

Community vote distribution
B (88%) 13%

srek3502 1 year, 1 month ago
30 Secs
Troubleshooting Splunk Enterprise Pg 44
upvoted 1 times

srek3502 1 year, 1 month ago
Correction: Troubleshooting Splunk Enterprise Pg 99
upvoted 1 times

qtygbapjpesdayazko 1 year, 5 months ago
Selected Answer: B
30 secs
upvoted 1 times

olibee 2 years, 2 months ago
First, metrics.log is a periodic report, taken every 30 seconds or so, of recent Splunk software activity. As written in the actual splunk documentation https://docs.splunk.com/Documentation/Splunk/latest/Troubleshooting/Aboutmetricslog
upvoted 2 times

RedYeti 2 years, 7 months ago
Selected Answer: B
30 seconds
https://docs.splunk.com/Documentation/Splunk/latest/Troubleshooting/Aboutmetricslog
upvoted 2 times

DimitarTsvetkov 2 years, 9 months ago
Selected Answer: B
As per PDF it's 30 secs
upvoted 4 times

SasnycoN 2 years, 10 months ago
Selected Answer: C
Answer "C" as per the troubleshooting PDF slide 106.
upvoted 1 times

IGoddard90 2 years, 9 months ago
troubleshooting PDF slide 106 states 30 seconds. So the answer is B.
upvoted 2 times

SasnycoN 2 years, 9 months ago
My bad. I meant Answer B - 30 seconds.
upvoted 2 times

Question #36 *Topic 1*

Which of the following is a good practice for a search head cluster deployer?

A. The deployer only distributes configurations to search head cluster members when they "phone home".

B. The deployer must be used to distribute non-replicable configurations to search head cluster members.

C. The deployer must distribute configurations to search head cluster members to be valid configurations.

D. The deployer only distributes configurations to search head cluster members with splunk apply shcluster-bundle.

Suggested Answer: *A*

Community vote distribution

B (100%)

sadhka Highly Voted 4 years, 2 months ago

My answer is B

upvoted 12 times

belindo96 4 years ago

Absolutely, answer A it's not a good practice its a fact.

upvoted 2 times

RedYeti 2 years, 7 months ago

it is not that it is a bad practice, search heads don't phone deployer...

upvoted 1 times

demarko 3 years, 11 months ago

B - https://docs.splunk.com/Documentation/Splunk/8.1.0/DistSearch/PropagateSHCconfigurationchanges

upvoted 3 times

lmak Most Recent 10 months, 1 week ago

The best practice for a search head cluster deployer is: D. The deployer only distributes configurations to search head cluster members with splunk apply shcluster-bundle.

A, B, and C, do not accurately reflect best practices:

A: Configurations should be pushed by the deployer actively rather than waiting for nodes to "phone home." This ensures timely and consistent distribution.

B: The deployer should be used for replicable configurations. Non-replicable configurations, such as user-specific configurations, should be managed locally on each search head.

C: The deployer should distribute configurations to ensure they are valid across the entire cluster. Waiting for configurations to become "valid" based on some criteria is not standard practice.

upvoted 2 times

srek3502 1 year, 1 month ago

Answer: B

https://docs.splunk.com/Documentation/Splunk/9.1.1/DistSearch/PropagateSHCconfigurationchanges

The deployer has these main roles:

It handles migration of app and user configurations into the search head cluster from non-cluster instances and search head pools.

It deploys baseline app configurations to search head cluster members.

It provides the means to distribute non-replicated, non-runtime configuration updates to all search head cluster members.

upvoted 1 times

esdee3 1 year, 2 months ago

I don't understand why the answer is A. Why "Phone Home"?

upvoted 1 times

qtygbapjpesdayazko 1 year, 5 months ago

Selected Answer: B

B. The deployer must be used to distribute non-replicable configurations to search head cluster members.

upvoted 1 times

Shahrukh_Salahudeen 1 year, 10 months ago

Selected Answer: B

Answer is B.

Explanation:

Option A: No, There is no phone home concept with regards to SearchHead DEPLOYER. phone_home is with Search-Head DEPLOYMENT SERVER.

Option B: Reference link - https://docs.splunk.com/Documentation/Splunk/9.0.3/DistSearch/PropagateSHCconfigurationchanges

It says:

"It provides the means to distribute non-replicated, non-runtime configuration updates to all search head cluster members."

Option C: No, deployer must NOT be used to distribute configurations to SH cluster members.

Option D: The statement is correct, but does not qualify to be a good practice.

upvoted 3 times

Redtonyeah 2 years, 8 months ago

Selected Answer: B

answer is B

upvoted 1 times

Crash_Override 2 years, 9 months ago

Selected Answer: B

The answer is B, A does not make sense because why would the deployer wait for a phone home when it uses the push mode to push apps to members.

upvoted 1 times

Question #37 *Topic 1*

A new Splunk customer is using syslog to collect data from their network devices on port 514. What is the best practice for ingesting this data into Splunk?

A. Configure syslog to send the data to multiple Splunk indexers.

B. Use a Splunk indexer to collect a network input on port 514 directly.

C. Use a Splunk forwarder to collect the input on port 514 and forward the data.

D. Configure syslog to write logs and use a Splunk forwarder to collect the logs.

Suggested Answer: *D*
Reference:
https://docs.splunk.com/Documentation/SplunkCloud/8.0.0/Data/Monitornetworkports

Community vote distribution
D (100%)

qtygbapjpesdayazko 1 year, 5 months ago

Selected Answer: D

D. Configure syslog to write logs and use a Splunk forwarder to collect the logs.
upvoted 1 times

Vale5M 1 year, 8 months ago

Answer is D. Show Data Admin slide 147
upvoted 2 times

RedYeti 2 years, 7 months ago

Selected Answer: D

Answer D
upvoted 2 times

david88f 3 years ago

Answer is: D
upvoted 3 times

demarko 3 years, 11 months ago

https://wiki.splunk.com/Community:BestPracticeForConfiguringSyslogInput
upvoted 2 times

Question #38 *Topic 1*

Which Splunk internal index contains license-related events?

A. _audit

B. _license

C. _internal

D. _introspection

Suggested Answer: *C*

Reference:

https://answers.splunk.com/answers/579494/how-to-display-license-consumed-by-an-index-over-2.html

Community vote distribution

C (100%)

qtygbapjpesdayazko 1 year, 5 months ago

Selected Answer: C

A. _audit, is for audit

B. _license, do not exist

C. _internal, have license data

D. _introspection, is use for splunk performance diag.

upvoted 2 times

minombrerodrigo 1 year, 10 months ago

Selected Answer: C

C is correct

upvoted 1 times

RedYeti 2 years, 7 months ago

Selected Answer: C

Answer C

upvoted 2 times

Bianchi 2 years, 7 months ago

Selected Answer: C

https://docs.splunk.com/Documentation/Splunk/8.2.6/Admin/LicenseUsageReportViewexamples

upvoted 1 times

just4learn 2 years, 8 months ago

_internal

upvoted 1 times

demarko 3 years, 11 months ago

https://docs.splunk.com/Documentation/Splunk/8.1.0/Admin/LicenseUsageReportViewexamples

upvoted 2 times

Question #39 *Topic 1*

Which of the following statements describe a Search Head Cluster (SHC) captain? (Select all that apply.)

A. Is the job scheduler for the entire SHC.

B. Manages alert action suppressions (throttling).

C. Synchronizes the member list with the KV store primary.

D. Replicates the SHC's knowledge bundle to the search peers.

Suggested Answer: *AD*
Reference:
https://docs.splunk.com/Documentation/Splunk/7.3.2/DistSearch/SHCarchitecture#role_of_the_captain

Community vote distribution
AB (100%)

ChantreyC Highly Voted 3 years, 10 months ago
A,B,D
https://docs.splunk.com/Documentation/Splunk/7.3.2/DistSearch/SHCarchitecture#role_of_the_captain
upvoted 8 times

SpTester 3 years, 5 months ago
Correct:
KV store can reside on a search head cluster. However, the search head cluster does not coordinate replication of KV store data or otherwise involve itself in the operation of KV store. For information on KV store, see About KV store in the Admin Manual.
upvoted 3 times

SpTester 3 years, 5 months ago
Actually yes.. page 219 of Cluster Admin pdf is correct too. it says: SHC captain and KV store primary synchronize their member lust everytime there is a status change -> so must be all ABCD than.
upvoted 6 times

qtygbapjpesdayazko 1 year, 3 months ago
ABCD
This is the way
upvoted 1 times

hamzaissam92 Highly Voted 1 year, 9 months ago
can't believe no one noticed this lol..D is definitely wrong, because the term "search peers" is another name for indexers and a captain doesn't replicate to search peers, it replicates to search head cluster members
upvoted 7 times

CactiAZ Most Recent 1 month ago
Answer is A,B,D

upvoted 1 times

Mntman77 6 months, 1 week ago
ABD

The KV store and search head clustering operate pretty much independently.

The KV store captain is elected by its peers, just like the search head captain, although it is different code that runs those two elections. It's possible that one node would end up hosting both captains, but not necessary.

The potential warning you point out is not a warning, but rather information about how kv store handles reads and writes. And where exactly do you think the docs are in conflict? KV store runs on a search head cluster, but search head clustering doesn't interact with the KV store.
upvoted 1 times

b5white 1 year, 4 months ago

Seems like C is wrong, per the comment that SpTester referenced, and the captain doesn't do anything with the peers, so AB.
upvoted 1 times

MaryKey 1 year, 4 months ago
Cannot select all options but all are valid here (ABCD).
upvoted 2 times

qtygbapjpesdayazko 1 year, 5 months ago
Selected Answer: AB
Is ABCD
upvoted 2 times

deepali_2710 1 year, 7 months ago
Option D is incorrect because the SHC captain is not responsible for replicating the SHC's knowledge bundle to the search peers. The SHC captain is responsible for managing the configuration and state of the SHC, including distributing the knowledge bundle to the other search heads in the cluster. However, it is the search peers that perform the actual searching against the indexed data, and they do not rely on the SHC captain to replicate the knowledge bundle.
upvoted 2 times

Hemnaath 1 year, 8 months ago
Answer is A, B & C
upvoted 2 times

Proctor 2 years, 1 month ago
ABCD

ABD: https://docs.splunk.com/Documentation/Splunk/9.0.1/DistSearch/SHCarchitecture#Search_head_cluster_captain

C: https://docs.splunk.com/Documentation/Splunk/9.0.1/Admin/ResyncKVstore#Resync_stale_KV_store_members
upvoted 5 times

qtygbapjpesdayazko 1 year, 3 months ago
This is the way.
upvoted 1 times

orezaie 2 years, 2 months ago
ABCD
C: Slide219 (cluster) SHC captain and KV store primary synchronize their member lust everytime there is a status change The captain also coordinates activities among all cluster members
D:slide186 cluster: Bundle replication to search peers happens only from the captain
upvoted 3 times

RedYeti 2 years, 7 months ago
Answers A, B and D

https://docs.splunk.com/Documentation/Splunk/latest/DistSearch/SHCarchitecture#Role_of_the_captain

In the last topic ("Search head clustering and KV store") we can read:
KV store can reside on a search head cluster. However, the search head cluster does not coordinate replication of KV store data or otherwise involve itself in the operation of KV store. For information on KV store, see About KV store in the Admin Manual.
upvoted 4 times

manu78 3 years, 7 months ago
>ABCD is correct
upvoted 3 times

[Removed] 3 years, 8 months ago
ABCD:

c: Slide219 (cluster) SHC captain and KV store primary synchronize their member lust everytime there is a status change

The captain also coordinates activities among all cluster members. Its responsibilities include:

Scheduling jobs. It assigns jobs to members, including itself, based on relative current loads.

Coordinating alerts and alert suppressions across the cluster. The captain tracks each alert but the member running an initiating search fires it.
Pushing the knowledge bundle to search peers.
Coordinating artifact replication. The captain ensures that search artifacts get replicated as necessary to fulfill the replication factor. See Choose the replication factor for the search head cluster.
Replicating configuration updates. The captain replicates any runtime changes to knowledge objects on one cluster member to all other members. This includes, for example, changes or additions to saved searches, lookup tables, and dashboards. See Configuration updates that the cluster replicates.

upvoted 3 times

[Removed] 3 years, 8 months ago

I think ABC is correct , D is not correct as it pushes the knowledge bundles to search peers

upvoted 2 times

Hamiltonian 3 years, 4 months ago

It's slippery terminology. In the pdf they say replicate, in the docs push.

upvoted 1 times

Question #40 *Topic 1*

Before users can use a KV store, an admin must create a collection. Where is a collection is defined?

A. kvstore.conf

B. collection.conf

C. collections.conf

D. kvcollections.conf

Suggested Answer: *C*

Reference:

https://docs.splunk.com/Documentation/Splunk/7.3.2/Knowledge/DefineaKVStorelookupinSplunkWeb

Community vote distribution

C (100%)

qtygbapjpesdayazko 1 year, 5 months ago

Selected Answer: C

Before you create a KV Store lookup, your Splunk deployment must have at least one KV Store collection. Certain apps, such as Enterprise Security, include KV Store collections with their installation.

KV Store collections are databases. They store your data as key/value pairs. When you create a KV Store lookup, the collection should have at least two fields. One of those fields should have a set of values that match with the values of a field in your event data, so that lookup matching can take place.

upvoted 1 times

RedYeti 2 years, 7 months ago

Selected Answer: C

Answer C

upvoted 3 times

asashima 2 years, 11 months ago

Answer is C.

See Clustering slides 217.

upvoted 4 times

david88f 3 years ago

collections.conf

upvoted 2 times

demarko 3 years, 11 months ago

https://dev.splunk.com/enterprise/docs/developapps/manageknowledge/kvstore/usingconfigurationfiles/

upvoted 2 times

Question #41 *Topic 1*

Which search will show all deployment client messages from the client (UF)?

A. index=_audit component=DC* host=<ds> | stats count by message

B. index=_audit component=DC* host=<uf> | stats count by message

C. index=_internal component= DC* host=<uf> | stats count by message

D. index=_internal component=DS* host=<ds> | stats count by message

Suggested Answer: *C*

Community vote distribution

C (100%)

SasnycoN Highly Voted 2 years, 10 months ago

Selected Answer: C

Answer C as per slide 83 in the Troubleshooting PDF

upvoted 6 times

KiranVM Most Recent 1 year, 8 months ago

Selected Answer: C

Answer C

upvoted 1 times

RedYeti 2 years, 7 months ago

Selected Answer: C

Answer C

Page 83 of Troubleshooting course

upvoted 2 times