Question #1 *Topic 1*

A global enterprise is running SAP ERP Central Component (SAP ECC) workloads on Oracle in an on-premises environment. The enterprise plans to migrate to SAP S/4HANA on AWS.
The enterprise recently acquired two other companies. One of the acquired companies is running SAP ECC on Oracle as its ERP system. The other acquired company is running an ERP system that is not from SAP. The enterprise wants to consolidate the three ERP systems into one ERP system on SAP S/4HANA on AWS. Not all the data from the acquired companies needs to be migrated to the final ERP system. The enterprise needs to complete this migration with a solution that minimizes cost and maximizes operational efficiency.
Which solution will meet these requirements?

A. Perform a lift-and-shift migration of all the systems to AWS. Migrate the ERP system that is not from SAP to SAP ECC. Convert all three systems to SAP S/4HANA by using SAP Software Update Manager (SUM) Database Migration Option (DMO). Consolidate all three SAP S/4HANA systems into a final SAP S/4HANA system. Decommission the other systems.

B. Perform a lift-and-shift migration of all the systems to AWS. Migrate the enterprise's initial system to SAP HANA, and then perform a conversion to SAP S/4HANA. Consolidate the two systems from the acquired companies with this SAP S/4HANA system by using the Selective Data Transition approach with SAP Data Management and Landscape Transformation (DMLT).

C. Use SAP Software Update Manager (SUM) Database Migration Option (DMO) with System Move to re-architect the enterprise's initial system to SAP S/4HANA and to change the platform to AWS. Consolidate the two systems from the acquired companies with this SAP S/4HANA system by using the Selective Data Transition approach with SAP Data Management and Landscape Transformation (DMLT).

D. Use SAP Software Update Manager (SUM) Database Migration Option (DMO) with System Move to re-architect all the systems to SAP S/4HANA and to change the platform to AWS. Consolidate all three SAP S/4HANA systems into a final SAP S/4HANA system. Decommission the other systems.

**Suggested Answer:** *A*

*Community vote distribution*

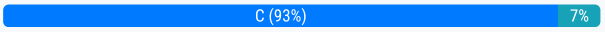


**chamorro** 1 year, 2 months ago
Did anyone take the exam recently? what is the % of accuracy?
upvoted 1 times

**XavilniestA** 1 year, 3 months ago
B and C omit to reduce the original SAP systems data to a selective subset, so it should be A or D.
In D I'm missing a dedicated handling of the non-SAP ERP,
so voting for A
upvoted 1 times

**Passexam4sure_com** 1 year, 8 months ago
Selected Answer: C
Use SAP Software Update Manager (SUM) Database Migration Option (DMO) with System Move to re-architect the enterprise's initial system to SAP S/4HANA and to change the platform to AWS. Consolidate the two systems from the acquired companies with this SAP S/4HANA system by using the Selective Data Transition approach with SAP Data Management and Landscape Transformation (DMLT).
upvoted 2 times

**vestersly** 1 year, 8 months ago
I vote D
upvoted 1 times

**forever_studious** 1 year, 11 months ago
A and D are wrong as one of the system is non-sap and DMP can not be used.
B gives the expected end result but C is more cost effective and efficient approach.
upvoted 2 times

**[Removed]** 1 year, 11 months ago
Selected Answer: B
Voting for B
upvoted 1 times

**[Removed]** 1 year, 11 months ago

EDIT C

upvoted 1 times

**ranac** 2 years, 2 months ago

**Selected Answer: C**

This option minimizes cost and maximizes operational efficiency by using the DMO with System Move to migrate the initial system to SAP S/4HANA on AWS. The Selective Data Transition approach with DMLT allows for the consolidation of the two acquired companies' systems with the new SAP S/4HANA system, enabling the migration of only necessary data and reducing overall complexity.

upvoted 2 times

**ranac** 2 years, 2 months ago

**Selected Answer: B**

This solution involves migrating all three systems to AWS, but only migrating the enterprise's initial system to SAP HANA and then converting it to SAP S/4HANA. The two acquired company systems will be consolidated with the SAP S/4HANA system by using the Selective Data Transition approach with SAP Data Management and Landscape Transformation (DMLT). This approach allows for selective data migration, which will help minimize costs and optimize operational efficiency.

upvoted 1 times

**ranac** 2 years, 2 months ago

Revisiting this again

C. Use SAP Software Update Manager (SUM) Database Migration Option (DMO) with System Move to re-architect the enterprise's initial system to SAP S/4HANA and to change the platform to AWS. Consolidate the two systems from the acquired companies with this SAP S/4HANA system by using the Selective Data Transition approach with SAP Data Management and Landscape Transformation (DMLT).

upvoted 1 times

**ranac** 2 years, 2 months ago

This option minimizes cost and maximizes operational efficiency by using the DMO with System Move to migrate the initial system to SAP S/4HANA on AWS. The Selective Data Transition approach with DMLT allows for the consolidation of the two acquired companies' systems with the new SAP S/4HANA system, enabling the migration of only necessary data and reducing overall complexity.

upvoted 1 times

**asharma45** 2 years, 4 months ago

surely C

upvoted 3 times

**ohcn** 2 years, 4 months ago

I think C

upvoted 2 times

**schalke04** 2 years, 4 months ago

**Selected Answer: C**

C is correct

upvoted 2 times

**Hyperdanny** 2 years, 4 months ago

C, since we don't have to move all the data.

upvoted 2 times

**kk8s** 2 years, 5 months ago

D, i think

upvoted 1 times

**dvrao** 2 years, 5 months ago

**Selected Answer: C**

yes, its C. because the migration should be cost effective as well

upvoted 2 times

**SMALLAM** 2 years, 5 months ago

I think

It's C

upvoted 4 times

Question #2 *Topic 1*

A global retail company is running its SAP landscape on AWS. Recently, the company made changes to its SAP Web Dispatcher architecture. The company added an additional SAP Web Dispatcher for high availability with an Application Load Balancer (ALB) to balance the load between the two SAP Web Dispatchers.
When users try to access SAP through the ALB, the system is reachable. However, the SAP backend system is showing an error message. An investigation reveals that the issue is related to SAP session handling and distribution of requests. The company confirmed that the system was working as expected with one SAP Web Dispatcher. The company replicated the configuration of that SAP Web Dispatcher to the new SAP Web Dispatcher.
How can the company resolve the error?

A. Maintain persistence by using session cookies. Enable session stickiness (session affinity) on the SAP Web Dispatchers by setting the wdisp/HTTP/esid_support parameter to True.

B. Maintain persistence by using session cookies. Enable session stickiness (session affinity) on the ALB.

C. Turn on host-based routing on the ALB to route traffic between the SAP Web Dispatchers.

D. Turn on URL-based routing on the ALB to route traffic to the application based on URL.

**Suggested Answer:** *C*

*Community vote distribution*


**Hyperdanny** Highly Voted 2 years, 4 months ago
B: "The company confirmed that the system was working as expected with one SAP Web Dispatcher." , so Web Dispatcher settings aren't the problem.
upvoted 11 times

**sagsgg** Highly Voted 2 years, 5 months ago
I think B
upvoted 8 times

**odre90** Most Recent 1 year, 6 months ago
Selected Answer: B
Must be B, the setting on webdispatcher affects how requests are distributed among the application servers - if the issue wasn't existing before, it must be connected to the multiple dispatchers setup. We must to make sure that once the user is connected to given dispatcher, it will remain so. B is the best answer then.
upvoted 2 times

**prisum** 1 year, 8 months ago
Option A seems to be correct.
Maintaining persistence through session cookies and enabling session stickiness on the SAP Web Dispatchers seems like a reasonable approach to ensure that user sessions are consistently directed to the same backend server. The parameter wdisp/HTTP/esid_support set to True should help in achieving this.
upvoted 1 times

**ggrodskiy** 1 year, 10 months ago
Correct B.
upvoted 1 times

**[Removed]** 1 year, 11 months ago
Selected Answer: B
Voting B
upvoted 2 times

**luckdragon** 1 year, 11 months ago
Selected Answer: B
B. the ALB is splitting the traffic, so stickiness needs to be set in the ALB, setting it in the dispatcher seems pointless of the ALB is still splitting the traffic between different dispatchers...
upvoted 3 times

**juanvepe** 1 year, 11 months ago

A.
The error message being displayed is related to SAP session handling and distribution of requests. By using session cookies, the company can maintain persistence of the user's session across requests. By enabling session stickiness on the SAP Web Dispatchers by setting the wdisp/HTTP/esid_support parameter to True, the company can ensure that requests from the same user are always routed to the same SAP Web Dispatcher.
This would resolve the error message that the company is seeing and ensure that the backend system is working as expected with the new SAP Web Dispatcher configuration.
upvoted 5 times

**anttan** 2 years, 4 months ago
A is correct.

To resolve the error related to SAP session handling and distribution of requests when using an Application Load Balancer (ALB) to balance the load between two SAP Web Dispatchers, the company should maintain session stickiness (session affinity) by using session cookies. This will ensure that each user's session is directed to the same SAP Web Dispatcher that the initial request was directed to.

Option A is the correct solution. The wdisp/HTTP/esid_support parameter is used to enable session stickiness on the SAP Web Dispatchers. By setting it to True, session information is maintained by the SAP Web Dispatcher in a session cookie. This cookie is then used to direct subsequent requests from the same user to the same SAP Web Dispatcher.

Option B is incorrect because enabling session stickiness on the ALB would not be enough to ensure that session information is maintained when requests are directed to different SAP Web Dispatchers.

Option C is incorrect because host-based routing would not address the issue of maintaining session information.

Option D is incorrect because URL-based routing is not designed to maintain session information.
upvoted 2 times

**ohcn** 2 years, 4 months ago
I think A. The issue is related to session stickness. Enabling wdisp/HTTP/esid_support you can maintain persistence of users session across requests.
upvoted 1 times

**schalke04** 2 years, 4 months ago
Selected Answer: A
A. https://help.sap.com/doc/saphelp_nw74/7.4.16/en-us/48/957c9194cc73eae10000000a42189b/content.htm?no_cache=true
upvoted 5 times

**SMALLAM** 2 years, 4 months ago
I think its A, as the issue is in coming from SAP side?
upvoted 2 times

**Hyperdanny** 2 years, 5 months ago
I am also voting for B. There is a connection between ALB and the dispatchers, but user sessions won't be handle correctly, therefore causing an error.
upvoted 3 times

Question #3 *Topic 1*

A company hosts its SAP NetWeaver workload on SAP HANA in the AWS Cloud. The SAP NetWeaver application is protected by a cluster solution that uses Red Hat Enterprise Linux. High Availability Add-On. The cluster solution uses an overlay IP address to ensure that the high availability cluster is still accessible during failover scenarios.
An SAP solutions architect needs to facilitate the network connection to this overlay IP address from multiple locations. These locations include more than 25 VPCs, other AWS Regions, and the on-premises environment. The company already has set up an AWS Direct Connect connection between the on-premises environment and AWS.
What should the SAP solutions architect do to meet these requirements in the MOST scalable manner?

A. Use VPC peering between the VPCs to route traffic between them.

B. Use AWS Transit Gateway to connect the VPCs and on-premises networks together.

C. Use a Network Load Balancer to route connections to various targets within VPCs.

D. Deploy a Direct Connect gateway to connect the Direct Connect connection over a private VIF to one or more VPCs in any accounts.

**Suggested Answer:** *D*

*Community vote distribution*

B (100%)

**venkyvenky890** 1 year, 7 months ago
I will also go with Option B.
1. It is not a transit, it is a private VIF which is provided in the answer
2. Private vif can connect only to 1 VPC at a time in a single region.
3. You would need 25 vifs to connect to 25 VPC's.
Hence Option B is correct
upvoted 2 times

**ggrodskiy** 1 year, 10 months ago
Correct B.
AWS Transit Gateway is a network transit hub that can interconnect thousands of VPCs and on-premises networks through a central gateway. This simplifies the network architecture and eliminates the need for complex peering relationships. AWS Transit Gateway also supports inter-Region peering, which enables the connection of transit gateways across different AWS Regions using the AWS global network. This way, the SAP NetWeaver workload on SAP HANA can be accessed from multiple locations with high performance and security.
upvoted 2 times

**[Removed]** 1 year, 11 months ago
Selected Answer: B
Voting B, many VCPs go with transit gateway
upvoted 1 times

**ohcn** 2 years, 4 months ago
I think B
upvoted 2 times

**schalke04** 2 years, 4 months ago
Selected Answer: B
The most scalable solution would be option B: Use AWS Transit Gateway to connect the VPCs and on-premises networks together.
upvoted 4 times

**SMALLAM** 2 years, 4 months ago
I am voting B
upvoted 2 times

**Hyperdanny** 2 years, 5 months ago
I am voting B
upvoted 2 times

**sagsgg** 2 years, 5 months ago
Selected Answer: B

I think B

upvoted 4 times

**kk8s** 2 years, 5 months ago

**Selected Answer: B**

B for me

upvoted 3 times

Question #4 *Topic 1*

A company is implementing SAP HANA on AWS. According to the company's security policy, SAP backups must be encrypted. Only authorized team members can have the ability to decrypt the SAP backups.
What is the MOST operationally efficient solution that meets these requirements?

A. Configure AWS Backint Agent for SAP HANA to create SAP backups in an Amazon S3 bucket. After a backup is created, encrypt the backup by using client-side encryption. Share the encryption key with authorized team members only.

B. Configure AWS Backint Agent for SAP HANA to use AWS Key Management Service (AWS KMS) for SAP backups. Create a key policy to grant decryption permission to authorized team members only.

C. Configure AWS Storage Gateway to transfer SAP backups from a file system to an Amazon S3 bucket. Use an S3 bucket policy to grant decryption permission to authorized team members only.

D. Configure AWS Backint Agent for SAP HANA to use AWS Key Management Service (AWS KMS) for SAP backups. Grant object ACL decryption permission to authorized team members only.

**Suggested Answer:** *C*

*Community vote distribution*

B (100%)

**schalke04** Highly Voted 2 years, 4 months ago

Selected Answer: B

The MOST operationally efficient solution that meets the requirements is option B: Configure AWS Backint Agent for SAP HANA to use AWS Key Management Service (AWS KMS) for SAP backups. Create a key policy to grant decryption permission to authorized team members only.

AWS KMS is a managed service that makes it easy to create and control encryption keys used to encrypt your data. By using AWS KMS to encrypt the backups, the encryption and decryption of the data is handled by AWS, freeing up the company's resources. Additionally, the key policy ensures that only authorized team members can decrypt the backups, thereby meeting the security requirements.

upvoted 9 times

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

B, since it is the most efficient option. A would work too, but requires more effort.

upvoted 5 times

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

Correct B.
AWS Backint Agent for SAP HANA to use AWS Key Management Service (AWS KMS) for SAP backups. AWS Backint Agent for SAP HANA is a tool that integrates SAP HANA with Amazon S3 and enables you to create and manage SAP HANA backups in Amazon S3https://docs.aws.amazon.com/sap/latest/sap-hana/aws-backint-agent-Amazon-S3.html. AWS KMS is a service that allows you to create and manage encryption keys and use them to encrypt and decrypt data in AWS services and in your applicationshttps://docs.aws.amazon.com/aws-backup/latest/devguide/encryption.html. By using AWS Backint Agent for SAP HANA with AWS KMS, you can encrypt your SAP backups with a customer master key (CMK) that you control and specify in the AWS Backup vault that stores your backupshttps://docs.aws.amazon.com/sap/latest/sap-hana/aws-backint-agent-Amazon-S3.html. You can also create a key policy to grant decryption permission to authorized team members only, which will ensure that only they can access and restore the encrypted backupshttps://docs.aws.amazon.com/aws-backup/latest/devguide/encryption.html.

upvoted 1 times

**kk8s** 2 years, 5 months ago

A
https://aws.amazon.com/blogs/apn/managing-sap-hana-database-backups-on-aws-with-syntax-cxlink-backup/

upvoted 1 times

**SMALLAM** 2 years, 5 months ago

I also think it's B

upvoted 4 times

**sagsgg** 2 years, 5 months ago

Selected Answer: B

I think B

upvoted 3 times

Question #5 *Topic 1*

A data analysis company has two SAP landscapes that consist of sandbox, development, QA, pre-production, and production servers. One landscape is on Windows, and the other landscape is on Red Hat Enterprise Linux. The servers reside in a room in a building that other tenants share.
An SAP solutions architect proposes to migrate the SAP applications to AWS. The SAP solutions architect wants to move the production backups to AWS and wants to make the backups highly available to restore in case of unavailability of an on-premises server.
Which solution will meet these requirements MOST cost-effectively?

A. Take a backup of the production servers. Implement an AWS Storage Gateway Volume Gateway. Create file shares by using the Storage Gateway Volume Gateway. Copy the backup files to the file shares through NFS and SMB.

B. Take a backup of the production servers. Send those backups to tape drives. Implement an AWS Storage Gateway Tape Gateway. Send the backups to Amazon S3 Standard-Infrequent Access (S3 Standard-IA) through the S3 console. Move the backups immediately to S3 Glacier Deep Archive.

C. Implement a third-party tool to take images of the SAP application servers and database server. Take regular snapshots at 1-hour intervals. Send the snapshots to Amazon S3 Glacier directly through the S3 Glacier console. Store the same images in different S3 buckets in different AWS Regions.

D. Take a backup of the production servers. Implement an Amazon S3 File Gateway. Create file shares by using the S3 File Gateway. Copy the backup files to the file shares through NFS and SMB. Map backup files directly to Amazon S3. Configure an S3 Lifecycle policy to send the backup files to S3 Glacier based on the company's data retention policy.

**Suggested Answer:** *C*

*Community vote distribution*

D (100%)

**schalke04** Highly Voted 2 years, 4 months ago

Selected Answer: D

This solution will allow the company to store their backups in Amazon S3, which is highly available and durable, and automate the movement of data to S3 Glacier based on the company's data retention policy, providing cost-effective data storage.

upvoted 5 times

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

Correct D.

upvoted 1 times

**SMALLAM** 2 years, 4 months ago

I think its D

upvoted 2 times

**Balki** 2 years, 4 months ago

Selected Answer: D

https://aws.amazon.com/blogs/storage/integrate-an-sap-ase-database-to-amazon-s3-using-aws-storage-gateway/

upvoted 4 times

**kk8s** 2 years, 5 months ago

D for me

https://aws.amazon.com/blogs/awsforsap/passive-disaster-recovery-for-sap-applications-using-aws-backup-and-aws-backint-agent/

upvoted 4 times

**sagsgg** 2 years, 5 months ago

Selected Answer: D

I think D

upvoted 3 times

Question #6 *Topic 1*

A company's SAP basis team is responsible for database backups in Amazon S3. The company frequently needs to restore the last 3 months of backups into the pre-production SAP system to perform tests and analyze performance. Previously, an employee accidentally deleted backup files from the S3 bucket. The SAP basis team wants to prevent accidental deletion of backup files in the future.
Which solution will meet these requirements?

A. Create a new resource-based policy that prevents deletion of the S3 bucket.

B. Enable versioning and multi-factor authentication (MFA) on the S3 bucket.

C. Create signed cookies for the backup files in the S3 bucket. Provide the signed cookies to authorized users only.

D. Apply an S3 Lifecycle policy to move the backup files immediately to S3 Glacier.

**Suggested Answer:** *A*

*Community vote distribution*

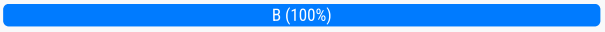


**Xavilniesta** 1 year, 3 months ago

B - MFA helps to avoid accidental deletion

upvoted 1 times

**Mojakerian** 1 year, 7 months ago

Correct is A

Every backup file is different in the bucket, there is no new versions of it. A policy for the whole bucket is much more effective and simple.

upvoted 3 times

**ggrodskiy** 1 year, 10 months ago

Correct B.

This option will allow the SAP basis team to enable versioning and multi-factor authentication (MFA) on the S3 bucket. Versioning is a feature that allows you to preserve, retrieve, and restore every version of every object stored in an S3 buckethttps://aws.amazon.com/getting-started/hands-on/protect-data-on-amazon-s3/. MFA is a security feature that requires users to provide two forms of authentication when performing certain actions on an S3 buckethttps://stackoverflow.com/questions/72634045/aws-s3-how-to-protect-against-accidental-deletion. By enabling versioning and MFA on the S3 bucket, the SAP basis team can protect their backup files from being overwritten or deleted by mistake or by unauthorized users. They can also recover any deleted versions of their backup files from the S3 bucket.

upvoted 3 times

**Dhieraj** 1 year, 10 months ago

Answer is B:

When versioning is enabled, a simple DELETE cannot permanently delete an object. Instead, Amazon S3 inserts a delete marker in the bucket, and that marker becomes the current version of the object with a new ID.

A wiil prevent all deletions..

upvoted 1 times

**schalke04** 2 years, 4 months ago

Selected Answer: B

B. Enable versioning and multi-factor authentication (MFA) on the S3 bucket.

upvoted 4 times

**SMALLAM** 2 years, 4 months ago

I think B

upvoted 2 times

**Hyperdanny** 2 years, 4 months ago

B, since inline policy (A) would prevent all deletes, not just accidental deletes.

upvoted 3 times

**bigmuramura2** 2 years, 5 months ago

Selected Answer: B

I think B

upvoted 3 times

**bigmuramura2** 2 years, 5 months ago

I think B

upvoted 2 times

Question #7 *Topic 1*

A company wants to run SAP HANA on AWS in the eu-central-1 Region. The company must make the SAP HANA system highly available by using SAP HANA system replication. In addition, the company must create a disaster recovery (DR) solution that uses SAP HANA system replication in the eu-west-1 Region. As prerequisites, the company has confirmed that Inter-AZ latency is less than 1 ms and that Inter-Region latency is greater than 1 ms.
Which solutions will meet these requirements? (Choose two.)

A. Install the tier 1 primary system and the tier 2 secondary system in eu-central-1. Configure the tier 1 system in Availability Zone 1. Configure the tier 2 system in Availability Zone 2. Configure SAP HANA system replication between tier 1 and tier 2 by using ASYNC replication mode. Install the DR tier 3 secondary system in eu-west-1 by using SYNC replication mode.

B. Install the tier 1 primary system and the tier 2 secondary system in eu-central-1. Configure the tier 1 system in Availability Zone 1. Configure the tier 2 system in Availability Zone 2. Configure SAP HANA system replication between tier 1 and tier 2 by using SYNC replication mode. Install the DR tier 3 secondary system in eu-west-1 by using ASYNC replication mode.

C. Install the tier 1 primary system and the tier 2 secondary system in eu-central-1. Configure the tier 1 system in Availability Zone 1. Configure the tier 2 system in Availability Zone 2. Configure SAP HANA system replication between tier 1 and tier 2 by using SYNC replication mode. Install the DR tier 3 secondary system in eu-west-1. Store daily backups from tier 1 in an Amazon S3 bucket in eu-central-1. Use S3 Cross-Region Replication to copy the daily backups to eu-west-1, where they can be restored if needed.

D. Install the tier 1 primary system in eu-central-1. Install the tier 2 secondary system and the DR tier 3 secondary system in eu-west-1. Configure the tier 2 system in Availability Zone 1. Configure the tier 3 system in Availability Zone 2. Configure SAP HANA system replication between all tiers by using ASYNC replication mode.

E. Install the tier 1 primary system and the tier 2 secondary system in eu-central-1. Configure the tier 1 system in Availability Zone 1. Configure the tier 2 system in Availability Zone 2. Configure SAP HANA system replication between tier 1 and tier 2 by using SYNCMEM replication mode. Install the DR tier 3 secondary system in eu-west-1 by using ASYNC replication mode.

**Suggested Answer:** *CD*

*Community vote distribution*


**Xavilniesta** 1 year, 3 months ago
I'd discard B only as SYNC replication is not apt for cross region with > 1 ms latency,
the other options look feasible
upvoted 1 times

**SAPEXAMS** 1 year, 7 months ago
**Selected Answer: BE**
B and E
upvoted 1 times

**Dhieraj** 1 year, 10 months ago
**Selected Answer: BE**
B and E are right answers.
upvoted 1 times

**[Removed]** 1 year, 11 months ago
**Selected Answer: BC**
Voting B,C
upvoted 1 times

**[Removed]** 1 year, 11 months ago
edit B,E
upvoted 1 times

**ohcn** 2 years, 4 months ago
**Selected Answer: BE**
B and E
upvoted 3 times

**schalke04** 2 years, 4 months ago

**Selected Answer: BE**

B E will work

upvoted 3 times

**SMALLAM** 2 years, 4 months ago

Agree with BE

upvoted 2 times

**kk8s** 2 years, 5 months ago

**Selected Answer: BE**

Agree with BE

https://docs.aws.amazon.com/sap/latest/sap-hana/hana-ops-patterns-multi.html

https://docs.aws.amazon.com/sap/latest/sap-hana/sap-hana-on-aws-aws-infrastructure-operating-system-setup-and-hana-installation.html

upvoted 2 times

**forexamweb** 2 years, 5 months ago

**Selected Answer: BE**

I think B and E.

The problem statement says that HANA system replication is used for both high availability and DR. Therefore, for High Availability, it would be SYNC replication mode or SYNCMEM replication mode, and for DR, it would be ASYNC replication mode. So, the answer is B and E.

Option C does not use HANA system replication in DR.

upvoted 4 times

**bigmuramura2** 2 years, 5 months ago

**Selected Answer: BC**

I think B and C

upvoted 2 times

Question #8 *Topic 1*

A company is running an SAP ERP Central Component (SAP ECC) system on an SAP HANA database that is 10 TB in size. The company is receiving notifications about long-running database backups every day. The company uses AWS Backint Agent for SAP HANA (AWS Backint agent) on an Amazon EC2 instance to back up the database. An SAP NetWeaver administrator needs to troubleshoot the problem and propose a solution. Which solution will help resolve this problem?

A. Ensure that AWS Backint agent is configured to send the backups to an Amazon S3 bucket over the internet. Ensure that the EC2 instance is configured to access the internet through a NAT gateway.

B. Check the UploadChannelSize parameter for AWS Backint agent. Increase this value in the aws-backint-agent-config.yaml configuration file based on the EC2 instance type and storage configurations.

C. Check the MaximumConcurrentFilesForRestore parameter for AWS Backint agent. Increase the parameter from 5 to 10 by using the aws-backint-agent-config.yaml configuration file.

D. Ensure that the backups are compressed. If necessary, configure AWS Backint agent to compress the backups and send them to an Amazon S3 bucket.

**Suggested Answer:** *A*

*Community vote distribution*

B (100%)

**juanvepe** 2 years ago

B:
https://docs.aws.amazon.com/sap/latest/sap-hana/aws-backint-agent-troubleshooting.html

Root Cause: The connection between AWS Backint agent and S3 fails due to high throughput.

Resolution: Use the following steps to troubleshoot this issue.

Update AWS Backint agent to version 1.02 or higher.

Lower the following backup and restore parameters:

Backup

UploadConcurrency

UploadChannelSize

Restore

MaximumConcurrentFilesForRestore

DownloadConcurrency

These values reduce concurrency and parallelism used by AWS Backint agent to achieve high performance during backup and restore.

Review network setup and configuration.

Perform trace route to see if Amazon S3 traffic goes through firewall package scanners or any other software that could significantly increase network latency.

upvoted 1 times

**schalke04** 2 years, 4 months ago

Selected Answer: B

B is fine.

upvoted 3 times

**SMALLAM** 2 years, 4 months ago

I think B

upvoted 2 times

**Hyperdanny** 2 years, 4 months ago

B:

The performance of backup and restore depends on many factors, such as the type of EC2 instance used, the EBS volumes, and the number of SAP HANA channels. If your database size is less than 128 GB, SAP HANA defaults to a single channel, or your SAP HANA parameter parallel_data_backup_backint_channels is set to 1.

upvoted 3 times

**Grillppl** 2 years, 4 months ago

B. Check the UploadChannelSize parameter

The UploadChannelSize parameter is used to determine how many files can be uploaded in parallel to the S3 bucket during backups.

https://docs.aws.amazon.com/sap/latest/sap-hana/aws-backint-agent-installing-configuring.html

upvoted 4 times

**kk8s** 2 years, 5 months ago

Maybe B

https://docs.aws.amazon.com/sap/latest/sap-hana/aws-backint-agent-troubleshooting.html

upvoted 2 times

**sagsgg** 2 years, 5 months ago

**Selected Answer: B**

I think B

upvoted 4 times

Question #9 *Topic 1*

A company wants to migrate its SAP workloads to AWS from another cloud provider. The company's landscape consists of SAP S/4HANA, SAP BW/4HANA, SAP Solution Manager, and SAP Web Dispatcher. SAP Solution Manager is running on SAP HANA.
The company wants to change the operating system from SUSE Linux Enterprise Server to Red Hat Enterprise Linux as a part of this migration.
The company needs a solution that results in the least possible downtime for the SAP S/4HANA and SAP BW/4HANA systems.
Which migration solution will meet these requirements?

A. Use SAP Software Provisioning Manager to perform a system export/import for SAP S/4HANA, SAP BW/4HANA, SAP Solution Manager, and SAP Web Dispatcher.

B. Use backup and restore for SAP S/4HANA, SAP BW/4HANA, and SAP Solution Manager. Reinstall SAP Web Dispatcher on AWS with the necessary configuration.

C. Use backup and restore for SAP S/4HANA and SAP BW/4HANA. Use SAP Software Provisioning Manager to perform a system export/import for SAP Solution Manager. Reinstall SAP Web Dispatcher on AWS with the necessary configuration.

D. Use SAP HANA system replication to replicate the data between the source system and the target AWS system for SAP S/4HANA and SAP BW/4HANA. Use SAP Software Provisioning Manager to perform a system export/import for SAP Solution Manager. Reinstall SAP Web Dispatcher on AWS with the necessary configuration.

**Suggested Answer:** *D*

*Community vote distribution*

D (100%)

**forexamweb** 2 years, 4 months ago

**Selected Answer: D**

The following REDHAT site describes the conversion of SAP HANA from SUSE to REDHAT.
SAP HANA system replication is used.
https://www.redhat.com/en/resources/migrating-sap-workloads-linux-detail

upvoted 4 times

**Hyperdanny** 2 years, 4 months ago

I still have doubts about D. Is it even possible to use System replication if the operating systems are different? Can someone open this :
https://userapps.support.sap.com/sap/support/knowledge/en/2763388

upvoted 1 times

**schalke04** 2 years, 4 months ago

**Selected Answer: D**

D. Use SAP HANA system replication to replicate the data between the source system and the target AWS system for SAP S/4HANA and SAP BW/4HANA. Use SAP Software Provisioning Manager to perform a system export/import for SAP Solution Manager. Reinstall SAP Web Dispatcher on AWS with the necessary configuration.

upvoted 3 times

**Hyperdanny** 2 years, 4 months ago

I am wondering if B might be more appropriate. I am not sure, if SAP Software Provisioning will be needed if Solutions Manager runs on HANA as well....

upvoted 1 times

**trashi** 2 years, 4 months ago

SolMan DB does not matter, question is about "least possible DT for S/4 and BW/4", therefore D shoud be correct

upvoted 2 times

**Hyperdanny** 2 years, 4 months ago

D
https://docs.aws.amazon.com/sap/latest/sap-hana/migrating-hana-hana-to-aws.html

upvoted 2 times

**kk8s** 2 years, 5 months ago

D for me.

https://docs.aws.amazon.com/migrationhub-orchestrator/latest/userguide/migrate-sap.html

upvoted 3 times

Question #10 *Topic 1*

A company is running an SAP on Oracle system on IBM Power architecture in an on-premises data center. The company wants to migrate the SAP system to AWS. The Oracle database is 15 TB in size. The company has set up a 100 Gbps AWS Direct Connect connection to AWS from the on-premises data center.
Which solution should the company use to migrate the SAP system MOST quickly?

A. Before the migration window, build a new installation of the SAP system on AWS by using SAP Software Provisioning Manager. During the migration window, export a copy of the SAP system and database by using the heterogeneous system copy process and R3load. Copy the output of the SAP system files to AWS through the Direct Connect connection. Import the SAP system to the new SAP installation on AWS. Switch over to the SAP system on AWS.

B. Before the migration window, build a new installation of the SAP system on AWS by using SAP Software Provisioning Manager. Back up the Oracle database by using native Oracle tools. Copy the backup of the Oracle database to AWS through the Direct Connect connection. Import the Oracle database to the SAP system on AWS. Configure Oracle Data Guard to begin replicating on-premises database log changes from the SAP system to the new AWS system. During the migration window, use Oracle to replicate any remaining changes to the Oracle database hosted on AWS. Switch over to the SAP system on AWS.

C. Before the migration window, build a new installation of the SAP system on AWS by using SAP Software Provisioning Manager. Create a staging Oracle database on premises to perform Cross Platform Transportable Tablespace (XTTS) conversion on the Oracle database. Take a backup of the converted staging database. Copy the converted backup to AWS through the Direct Connect connection. Import the Oracle database backup to the SAP system on AWS. Take regularly scheduled incremental backups and XTTS conversions of the staging database. Transfer these backups and conversions to the AWS target database. During the migration window, perform a final incremental Oracle backup. Convert the final Oracle backup by using XTTS. Replay the logs in the target Oracle database hosted on AWS. Switch over to the SAP system on AWS.

D. Before the migration window, launch an appropriately sized Amazon EC2 instance on AWS to receive the migrated SAP database. Create an AWS Server Migration Service (AWS SMS) job to take regular snapshots of the on-premises Oracle hosts. Use AWS SMS to copy the snapshot as an AMI to AWS through the Direct Connect connection. Create a new SAP on Oracle system by using the migrated AMI. During the migration window, take a final incremental SMS snapshot and copy the snapshot to AWS. Restart the SAP system by using the new up-to-date AMI. Switch over to the SAP system on AWS.

**Suggested Answer:** *A*

*Community vote distribution*

C (100%)

**kk8s** Highly Voted 2 years, 5 months ago

I think C.

https://aws.amazon.com/blogs/awsforsap/reducing-downtime-with-oracle-xtts-method-for-cross-platform-sap-migrations/

upvoted 9 times

**khchan123** Most Recent 1 year, 6 months ago

D is the quickest. Taking AMI snapshot and restore is usually the most simple and quickest way. Data transfer time (20 mins) is not the major component so compression may not benefit a lot.

upvoted 1 times

**khchan123** 1 year, 6 months ago

Correction - C is correct.

Migrating from IBM Power to AWS is heterogenous, and not supported by AWS MGN.

upvoted 1 times

**Riskate** 1 year, 11 months ago

Selected Answer: C

C i think

upvoted 1 times

**juanvepe** 1 year, 11 months ago

Answer C.

https://aws.amazon.com/es/blogs/awsforsap/reducing-downtime-with-oracle-xtts-method-for-cross-platform-sap-migrations/

upvoted 1 times

**schalke04** 2 years, 4 months ago

**Selected Answer: C**

C makes sense after reading kk8s

upvoted 4 times

Question #11 *Topic 1*

An SAP solutions architect is designing an SAP HANA scale-out architecture for SAP Business Warehouse (SAP BW) on SAP HANA on AWS. The SAP solutions architect identifies the design as a three-node scale-out deployment of xte.32xiarge Amazon EC2 instances.
The SAP solutions architect must ensure that the SAP HANA scale-out nodes can achieve the low-latency and high-throughput network performance that are necessary for node-to-node communication.
Which combination of steps should the SAP solutions architect take to meet these requirements? (Choose two.)

A. Create a cluster placement group. Launch the instances into the cluster placement group.

B. Create a spread placement group. Launch the instances into the spread placement group.

C. Create a partition placement group. Launch the instances into the partition placement group.

D. Based on the operating system version, verify that enhanced networking is enabled on all the nodes.

E. Switch to a different instance family that provides network throughput that is greater than 25 Gbps.

**Suggested Answer:** *AB*

*Community vote distribution*

AD (100%)

**[Removed]** 1 year, 11 months ago

**Selected Answer: AD**

Voting A,D

upvoted 1 times

**SMALLAM** 2 years, 4 months ago

I think A&D

upvoted 2 times

**schalke04** 2 years, 4 months ago

**Selected Answer: AD**

A and D:

upvoted 3 times

**kk8s** 2 years, 5 months ago

AE for me

upvoted 1 times

**kk8s** 2 years, 4 months ago

sorry. change to AD

upvoted 2 times

**sagsgg** 2 years, 5 months ago

**Selected Answer: AD**

I think A D

upvoted 4 times

Question #12 *Topic 1*

A company needs to migrate its critical SAP workloads from an on-premises data center to AWS. The company has a few source production databases that are 10 TB or more in size. The company wants to minimize the downtime for this migration.
As part of the proof of concept, the company used a low-speed, high-latency connection between its data center and AWS. During the actual migration, the company wants to maintain a consistent connection that delivers high bandwidth and low latency. The company also wants to add a layer of connectivity resiliency. The backup connectivity does not need to be as fast as the primary connectivity.
An SAP solutions architect needs to determine the optimal network configuration for data transfer. The solution must transfer the data with minimum latency.
Which configuration will meet these requirements?

A. Set up one AWS Direct Connect connection for connectivity between the on-premises data center and AWS. Add an AWS Site-to-Site VPN connection as a backup to the Direct Connect connection.

B. Set up an AWS Direct Connect gateway with multiple Direct Connect connections that use a link aggregation group (LAG) between the on-premises data center and AWS.

C. Set up Amazon Elastic File System (Amazon EFS) file system storage between the on-premises data center and AWS. Configure a cron job to copy the data into this EFS mount. Access the data in the EFS file system from the target environment.

D. Set up two redundant AWS Site-to-Site VPN connections for connectivity between the on-premises data center and AWS.

**Suggested Answer:** *D*

*Community vote distribution*

A (100%)

**sagsgg** Highly Voted 2 years, 5 months ago

Selected Answer: A

I think A

upvoted 9 times

**schalke04** Most Recent 2 years, 4 months ago

Selected Answer: A

A: one direct connect, one VPN

upvoted 4 times

**kk8s** 2 years, 5 months ago

Selected Answer: A

Yes. A

upvoted 3 times

Question #13 *Topic 1*

A company wants to migrate its SAP ERP landscape to AWS. The company will use a highly available distributed deployment for the new architecture. Clients will access SAP systems from a local data center through an AWS Site-to-Site VPN connection that is already in place. An SAP solutions architect needs to design the network access to the SAP production environment.
Which configuration approaches will meet these requirements? (Choose two.)

A. For the ASCS instance, configure an overlay IP address that is within the production VPC CIDR range. Create an AWS Transit Gateway. Attach the VPN to the transit gateway. Use the transit gateway to route the communications between the local data center and the production VPC. Create a static route on the production VPC to route traffic that is directed to the overlay IP address to the ASCS instance.

B. For the ASCS instance, configure an overlay IP address that is outside the production VPC CIDR range. Create an AWS Transit Gateway. Attach the VPN to the transit gateway. Use the transit gateway to route the communications between the local data center and the production VPC. Create a static route on the production VPC to route traffic that is directed to the overlay IP address to the ASCS instance.

C. For the ASCS instance, configure an overlay IP address that is within the production VPC CIDR range. Create a target group that points to the overlay IP address. Create a Network Load Balancer, and register the target group. Create a static route on the production VPC to route traffic that is directed to the overlay IP address to the ASCS instance.

D. For the ASCS instance, configure an overlay IP address that is outside the production VPC CIDR range. Create a target group that points to the overlay IP address. Create a Network Load Balancer, and register the target group. Create a static route on the production VPC to route traffic that is directed to the overlay IP address to the ASCS instance.

E. For the ASCS instance, configure an overlay IP address that is outside the production VPC CIDR range. Create a target group that points to the overlay IP address. Create an Application Load Balancer, and register the target group. Create a static route on the production VPC to route traffic that is directed to the overlay IP address to the ASCS instance.

**Suggested Answer:** *BE*

*Community vote distribution*

BD (92%) 8%

**One_picese** 1 year, 6 months ago

Answer：BD
SAP on AWS High Availability with Overlay IP Address Routing：
AWS Transit Gateway serves as central hub to facilitate network connection to an overlay IP address.
Elastic Load Balancing where a Network Load Balancer enables network access to an overlay IP addres
https://docs.aws.amazon.com/sap/latest/sap-hana/sap-ha-overlay-ip.html
upvoted 1 times

**SONALID** 1 year, 8 months ago

Application load balancer cannot be attached with Overlay IP
upvoted 1 times

**Dhieraj** 1 year, 10 months ago

Selected Answer: BD
Overlay IP should be outside Prod VPC CIDR, thus A and C are eliminated.
HA of ASCS uses Network load balancer and not application load balancer, so E is eliminated
upvoted 3 times

**[Removed]** 1 year, 11 months ago

Selected Answer: AC
voting A,C
upvoted 1 times

**[Removed]** 1 year, 11 months ago

Edit B,D
upvoted 1 times

**juanvepe** 1 year, 11 months ago

Answers.
B-D

upvoted 2 times

**easytoo** 2 years ago

a-c

should be within the prod VPC CIDR range.

upvoted 1 times

**easytoo** 2 years ago

corrected - b-d

upvoted 1 times

**Shaktimaan** 2 years, 3 months ago

Why application load balancer and not network load balancer ?

upvoted 1 times

**easytoo** 2 years ago

the Application Load Balancer is not a good choice for this application because it is designed for HTTP traffic, not SAP traffic

upvoted 1 times

**SMALLAM** 2 years, 4 months ago

I think BD

upvoted 3 times

**schalke04** 2 years, 4 months ago

**Selected Answer: BD**

B and D

upvoted 2 times

**chopperharu** 2 years, 5 months ago

**Selected Answer: BD**

I think B,D

upvoted 4 times

**lunachi4** 2 years, 5 months ago

**Selected Answer: BD**

I think B/D

upvoted 3 times

Question #14 *Topic 1*

A company is running an SAP HANA database on AWS. The company is running AWS Backint Agent for SAP HANA (AWS Backint agent) on an Amazon EC2 instance. AWS Backint agent is configured to back up to an Amazon S3 bucket. The backups are failing with an AccessDenied error in the AWS Backint agent log file.
What should an SAP basis administrator do to resolve this error?

A. Assign execute permissions at the operating system level for the AWS Backint agent binary and for AWS Backint agent.

B. Assign an IAM role to an EC2 instance. Attach a policy to the IAM role to grant access to the target S3 bucket.

C. Assign the correct Region ID for the S3BucketAwsRegion parameter in AWS Backint agent for the SAP HANA configuration file.

D. Assign the value for the EnableTagging parameter in AWS Backint agent for the SAP HANA configuration file.

**Suggested Answer:** *C*

*Community vote distribution*

B (100%)

**chopperharu** Highly Voted 2 years, 5 months ago

Selected Answer: B

I think B

upvoted 6 times

**prisum** Most Recent 1 year, 8 months ago

B is correct

upvoted 1 times

**luckdragon** 1 year, 11 months ago

Selected Answer: B

who decides the "correct answer"?? they are almost all wrong, this is B

upvoted 1 times

**easytoo** 2 years ago

b-b-b-b-b-b-b-

upvoted 1 times

**juanvepe** 2 years ago

B. https://docs.aws.amazon.com/sap/latest/sap-hana/aws-backint-agent-troubleshooting.html

Problem: AccessDenied appears in agent logs.
Root Causes:

The IAM role for the EC2 instance does not have the correct permissions to access the S3 bucket.

The agent configuration file does not have the S3BucketOwnerAccountID in double quotes. The S3BucketOwnerAccountID is the 12-digit AWS Account ID.

The S3 bucket is not owned by the provided account for S3BucketOwnerAccountID.

The S3 bucket provided for the S3BucketOwnerAccountID was created before May 2019.

Resolution: Verify the prerequisite steps for installing the AWS Backint agent.

upvoted 2 times

**schalke04** 2 years, 4 months ago

Selected Answer: B

B. Assign an IAM role to an EC2 instance. Attach a policy to the IAM role to grant access to the target S3 bucket

upvoted 3 times

**kk8s** 2 years, 5 months ago

Yeah. B

upvoted 2 times

Question #15 *Topic 1*

A company is starting a new project to implement an SAP landscape with multiple accounts that belong to multiple teams in the us-east-2 Region. These teams include procurement, finance, sales, and human resources. An SAP solutions architect has started designing this new landscape and the AWS account structures.
The company wants to use automation as much as possible. The company also wants to secure the environment, implement federated access to accounts, centralize logging, and establish cross-account security audits. In addition, the company's management team needs to receive a top-level summary of policies that are applied to the AWS accounts.
What should the SAP solutions architect do to meet these requirements?

A. Use AWS CloudFormation StackSets to apply SCPs to multiple accounts in multiple Regions. Use an Amazon CloudWatch dashboard to check the applied policies in the accounts.

B. Use an AWS Elastic Beanstalk blue/green deployment to create IAM policies and apply them to multiple accounts together. Use an Amazon CloudWatch dashboard to check the applied policies in the accounts.

C. Implement guardrails by using AWS CodeDeploy and AWS CodePipeline to deploy SCPs into each account. Use the CodePipeline deployment dashboard to check the applied policies in the accounts.

D. Apply SCPs through AWS Control Tower. Use the AWS Control Tower integrated dashboard to check the applied policies in the accounts.

**Suggested Answer:** *D*

*Community vote distribution*


**schalke04** Highly Voted 2 years, 4 months ago

Selected Answer: D

The SAP solutions architect should do D. Apply SCPs through AWS Control Tower. Use the AWS Control Tower integrated dashboard to check the applied policies in the accounts.

AWS Control Tower is a service that automates the set-up of a multi-account AWS environment and provides guardrails to help enforce compliance and security best practices. By using AWS Control Tower, the company can apply SCPs to multiple accounts, monitor policies across multiple accounts, and receive a top-level summary of policies that are applied to the AWS accounts. Additionally, the integrated dashboard of AWS Control Tower can be used to check the applied policies in the accounts.

upvoted 5 times

**easytoo** Most Recent 2 years ago

d-d-d-dd-d-

upvoted 2 times

**Balki** 2 years, 4 months ago

Selected Answer: D

https://aws.amazon.com/blogs/mt/managing-the-multi-account-environment-using-aws-organizations-and-aws-control-tower/

upvoted 3 times

Question #16 *Topic 1*

A company is running its SAP workloads on premises and needs to migrate the workloads to AWS. All the workloads are running on SUSE Linux Enterprise Server and Oracle Database. The company's landscape consists of SAP ERP Central Component (SAP ECC), SAP Business Warehouse (SAP BW), and SAP NetWeaver systems. The company has a dedicated AWS Direct Connect connection between its on-premises environment and AWS. The company needs to migrate the systems to AWS with the least possible downtime.
Which migration solution will meet these requirements?

A. Use SAP Software Provisioning Manager to perform an export of the systems. Copy the export to Amazon S3. Use SAP Software Provisioning Manager to perform an import of the systems to SUSE Linux Enterprise Server and Oracle Database on AWS.

B. Use SAP Software Provisioning Manager to perform parallel export/import of the systems to migrate the systems to SUSE Linux Enterprise Server and Oracle Database on AWS.

C. Use SAP Software Provisioning Manager to perform parallel export/import of the systems to migrate the systems to Oracle Enterprise Linux and Oracle Database on AWS.

D. Use SAP Software Provisioning Manager to perform an export of the systems. Copy the export to Amazon S3. Use SAP Software Provisioning Manager to perform an import of the systems to Oracle Enterprise Linux and Oracle Database on AWS.

**Suggested Answer:** *C*

*Community vote distribution*

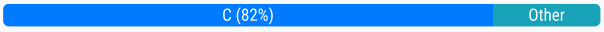


**Dhieraj** 1 year, 10 months ago

Selected Answer: C

SUSE does not support Oracle on AWS. So A and B are eliminated. Of C and D, parallel exp/imp will take less time so less outage

upvoted 3 times

**[Removed]** 1 year, 11 months ago

Selected Answer: A

going for A

https://sapblog.protiviti.com/2023/02/16/migrating-sap-workloads-to-aws-which-option-to-choose/

upvoted 1 times

**[Removed]** 1 year, 11 months ago

Sorry C

upvoted 1 times

**easytoo** 2 years ago

c-c-c-c-c-c

upvoted 2 times

**schalke04** 2 years, 4 months ago

Selected Answer: C

C . 100%

upvoted 3 times

**Hyperdanny** 2 years, 4 months ago

C: Only ORACLE Enterprise Linux supports SAP workloads on AWS

upvoted 3 times

**Kiran1982** 2 years, 4 months ago

Selected Answer: C

SUSE Linux is not supported by Oracle for SAP

upvoted 2 times

**MKM** 2 years, 4 months ago

It's C. Because Oracle DB can not run on SUSE Linux Enterprise Server on AWS. It can only run on Oracle Enterprise Linux on AWS.

upvoted 2 times

**forexamweb** 2 years, 4 months ago

Selected Answer: C

It is mandatory to have Oracle Enterprise Linux(OEL) as the operating system for running Oracle database for SAP.
https://launchpad.support.sap.com/#/notes/1656250

upvoted 1 times

**Balki** 2 years, 4 months ago

**Selected Answer: B**

parallel export and SUSE Linux Enterprise Server on Amazon

upvoted 1 times

Question #17 *Topic 1*

A company is designing a disaster recovery (DR) strategy for an SAP HANA database that runs on an Amazon EC2 instance in a single Availability Zone. The company can tolerate a long RTO and an RPO greater than zero if it means that the company can save money on its DR process.
The company has configured an Amazon CloudWatch alarm to automatically recover the EC2 instance if the instance experiences an unexpected issue. The company has set up AWS Backint Agent for SAP HANA to save the backups into Amazon S3.
What is the MOST cost-effective DR option for the company's SAP HANA database?

A. Set up AWS CloudFormation to automatically launch a new EC2 instance for the SAP HANA database in a second Availability Zone from backups that are stored in Amazon S3. When the SAP HANA database is operational, perform a database restore by using the standard SAP HANA restore process.

B. Launch a secondary EC2 instance for the SAP HANA database on a less powerful EC2 instance type in a second Availability Zone. Configure SAP HANA system replication with the preload option turned off.

C. Launch a secondary EC2 instance for the SAP HANA database on an equivalent EC2 instance type in a second Availability Zone. Configure SAP HANA system replication with the preload option turned on.

D. Set up AWS CloudFormation to automatically launch a new EC2 instance for the SAP HANA database in a second Availability Zone from backups that are stored in Amazon Elastic Block Store (Amazon EBS). When the SAP HANA database is operational, perform a database restore by using the standard SAP HANA restore process.

**Suggested Answer:** *A*

*Community vote distribution*


**schalke04** Highly Voted 2 years, 4 months ago

Selected Answer: A

A. restore from S3 is the cheapest

upvoted 6 times

**Kiran1982** Highly Voted 2 years, 4 months ago

Selected Answer: A

A is the most cost effective DR option.

upvoted 5 times

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

Selected Answer: A

A is the most cost effective since customer does not have specific requirements for RTO/RPO.

B is a warm standby with no pre-load setup, it is definitely not the most cost effective

upvoted 2 times

**kaishin0527** 1 year, 11 months ago

Selected Answer: A

A: This option leverages the existing backups stored in Amazon S3 and does not require running additional, potentially costly, EC2 instances until needed. The new instance is only launched in the event of a disaster, and data is then restored from the S3 backup. This is a cost-effective method, as you only pay for the S3 storage of your backup and the EC2 resources when you need them.

upvoted 1 times

**[Removed]** 1 year, 11 months ago

Selected Answer: A

A, use s3

upvoted 1 times

**easytoo** 2 years ago

b-b-bb-b-b

upvoted 1 times

**easytoo** 2 years ago

Considering the cost-effectiveness and the company's willingness to tolerate a long RTO and an RPO greater than zero, option B is the most suitable choice. It allows for the launch of a secondary EC2 instance in a different Availability Zone with a less powerful instance type and system replication configured without the preload option, potentially reducing costs while still providing some level of disaster recovery

upvoted 2 times

**djwhowe** 2 years, 1 month ago

**Selected Answer: B**

Cost optimized

You can reduce costs by using a smaller or shared secondary SAP HANA system. In the smaller secondary option, the infrastructure is initially sized smaller than the primary and resized before performing a takeover. In the shared secondary option, the unused memory on the secondary system is used by a non-production or sacrificial instance.

upvoted 2 times

**Deepk12493** 2 years, 4 months ago

B is the right answer.

A is incorrect. If you are able to restore when the primary DB is operational it is not a disaster scenario.

upvoted 1 times

**forexamweb** 2 years, 4 months ago

**Selected Answer: A**

A maybe

https://docs.aws.amazon.com/sap/latest/sap-hana/hana-ops-ha-dr.html#backint-hana-hadr

upvoted 2 times

**Balki** 2 years, 4 months ago

**Selected Answer: B**

https://docs.aws.amazon.com/sap/latest/sap-hana/hana-ops-ha-dr-hsr.html

upvoted 1 times

Question #18 *Topic 1*

A company is using a multi-account strategy for SAP HANA and SAP BW/4HANA instances across development, QA, and production systems in the same AWS Region. Each system is hosted in its own VPC. The company needs to establish cross-VPC communication between the SAP systems.
The company might add more SAP systems in the future. The company must create connectivity across the SAP systems and hundreds of AWS accounts. The solution must maximize scalability and reliability.
Which solution will meet these requirements?

A. Create an AWS Transit Gateway in a central networking account. Attach the transit gateway to the AWS accounts. Set up routing and a network ACL to establish communication.

B. Set up VPC peering between the accounts. Configure routing in each VPC to use the VPC peering links.

C. Create a transit VPC that uses the hub-and-spoke model. Set up routing to use the transit VPC for communication between the SAP systems.

D. Create a VPC link for each SAP system. Use the VPC links to connect the SAP systems.

**Suggested Answer:** *B*

*Community vote distribution*

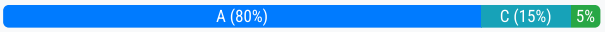


**awsmonster** 1 year, 4 months ago

**Selected Answer: A**

Answer is A.

https://aws.amazon.com/blogs/networking-and-content-delivery/migrate-from-transit-vpc-to-aws-transit-gateway/

As per the link, "Transit gateways are easy to set up and to use, and are designed to be highly scalable and resilient." While Transit-VPC "leverages instance-based routing that increases costs while lowering availability and limiting the bandwidth."

upvoted 3 times

**rrshah83** 1 year, 7 months ago

**Selected Answer: C**

The only reason A is incorrect because you can't attach transit gateway to accounts. You have vpc attachments. Hence chosing the next best option C: Transit vpc.

upvoted 2 times

**venkyvenky890** 1 year, 7 months ago

But you can share Transit gateways across AWS Accounts using RAM which can can then be used to route traffic between VPC's

upvoted 1 times

**kaishin0527** 1 year, 11 months ago

**Selected Answer: A**

A: AWS Transit Gateway is a service that enables you to connect your Amazon Virtual Private Clouds (VPCs) and on-premises networks to a single gateway. It is a great solution for scenarios like this where you need to connect multiple VPCs across several AWS accounts. It simplifies network architecture, reduces operational overhead, and is highly scalable. With Transit Gateway, you can set up routing and network ACLs to establish and control communication between the connected VPCs.

upvoted 1 times

**[Removed]** 1 year, 11 months ago

**Selected Answer: A**

A, transit gateway

upvoted 1 times

**easytoo** 2 years ago

a-a-aa-a-

upvoted 1 times

**schalke04** 2 years, 4 months ago

Selected Answer: A

A: The AWS Transit Gateway provides a scalable and centralized solution for cross-VPC communication and can support hundreds of AWS accounts. It allows creating a central hub that enables communication between multiple VPCs, providing a scalable and reliable solution for cross-VPC communication, which meets the company's requirements.

upvoted 4 times

**Kiran1982** 2 years, 4 months ago

Selected Answer: A

Transit gateway is best way to connect multiple VPC

upvoted 3 times

**Balki** 2 years, 4 months ago

Selected Answer: A

Hub and spoke. 100+ accounts. Transit gateway. Just cleared Networking specialty exam

upvoted 4 times

**kk8s** 2 years, 5 months ago

Selected Answer: C

C

https://aws.amazon.com/blogs/networking-and-content-delivery/creating-a-single-internet-exit-point-from-multiple-vpcs-using-aws-transit-gateway/

https://docs.aws.amazon.com/sap/latest/sap-hana/sap-oip-architecture.html

upvoted 1 times

**chopperharu** 2 years, 5 months ago

Selected Answer: B

I think B

upvoted 1 times

**forexamweb** 2 years, 5 months ago

I think A.

upvoted 4 times

Question #19 *Topic 1*

A company is planning to deploy a new SAP NetWeaver ABAP system on AWS with an Oracle database that runs on an Amazon EC2 instance. The EC2 instance uses a Linux-based operating system. The company needs a database storage solution that provides flexibility to adjust the IOPS regardless of the allocated storage size.
Which solution will meet these requirements MOST cost-effectively?

A. General Purpose SSD (gp3) Amazon Elastic Block Store (Amazon EBS) volumes

B. Amazon Elastic File System (Amazon EFS) Standard-Iinfrequent Access (Standard-IA) storage class

C. Amazon FSx for Windows File Server

D. Provisioned IOPS SSD (io2) Amazon Elastic Block Store (Amazon EBS) volumes

**Suggested Answer:** *A*

*Community vote distribution*

A (100%)

**kk8s** Highly Voted 2 years, 5 months ago

Selected Answer: A

A. both io2 and gp3 can resize IOPS
but if the cost effective, gp3 is the one.

https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/requesting-ebs-volume-modifications.html

https://aws.amazon.com/blogs/storage/optimizing-sap-hanas-persistence-layer-with-amazon-ebs-gp3-volumes/

upvoted 5 times

**kaishin0527** Most Recent 1 year, 11 months ago

Selected Answer: A

A: Amazon EBS gp3 volumes are the next generation of general purpose SSD volumes for Amazon EC2. gp3 volumes provide the ability to independently configure volume capacity and performance, enabling you to optimize costs for a wide range of workloads.

upvoted 3 times

**[Removed]** 1 year, 11 months ago

Selected Answer: A

gp3 is recommended

upvoted 1 times

**easytoo** 2 years ago

d-d-d-dd-

upvoted 1 times

**easytoo** 2 years ago

most cost effective is a-a-a-a-a

upvoted 1 times

**Balki** 2 years, 4 months ago

Selected Answer: A

gp2 and gp3 volumes balance price and performance for a variety of workloads, while io1, io2, and io2 Block Express volumes provide the highest performance for mission-critical applications. From these options, you can choose the best storage solution that meets your performance and cost requirements. We recommend the io2 or io2 Block Express configuration for mission-critical SAP HANA production workloads.

upvoted 3 times

**schalke04** 2 years, 4 months ago

Selected Answer: A

use GP3 by default

upvoted 3 times

Question #20 *Topic 1*

A company is using SAP NetWeaver with Java on AWS. The company has updated its generation of Amazon EC2 instances to the most recent generation of EC2 instances. When the company tries to start SAP, the startup fails. The log indicates that the SAP license expired or is not valid. What is the reason for this issue?

A. The instance ID changed as part of the EC2 generation change.

B. The instance's hypervisor changed from Xen to Nitro.

C. The SAP Java Virtual Machine (SAP JVM) is not compatible with the new instance type.

D. An EC2 generation change is not supported for SAP Java-based systems.

**Suggested Answer:** *A*

*Community vote distribution*

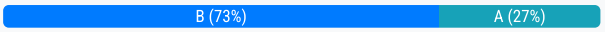


**anttan** Highly Voted 2 years, 4 months ago

Answer is B.

This change in hypervisor can cause the MAC address of the network interface to change, which in turn can cause the SAP license to be invalid. This is because the SAP license is tied to the MAC address of the network interface. To resolve the issue, the company should contact SAP to obtain a new license key that is valid for the new MAC address.

upvoted 6 times

**schalke04** Highly Voted 2 years, 4 months ago

Selected Answer: B

https://launchpad.support.sap.com/#/notes/2591601

upvoted 5 times

**prisum** Most Recent 1 year, 8 months ago

Option B is the irght answer

upvoted 1 times

**kaishin0527** 1 year, 11 months ago

Selected Answer: B

B: When you change the generation of the Amazon EC2 instance, the underlying hypervisor can change from Xen to Nitro. The SAP license is bound to the hypervisor's hardware key, so if the hypervisor changes, the hardware key changes, which can make the SAP license invalid. This is likely the reason why the startup is failing after the update to a newer generation of EC2 instances.

upvoted 1 times

**[Removed]** 1 year, 11 months ago

Selected Answer: A

https://docs.aws.amazon.com/sap/latest/general/overview-sap-on-aws.html

upvoted 1 times

**SMALLAM** 2 years, 4 months ago

I think B

upvoted 2 times

**matakuyy2** 2 years, 4 months ago

Selected Answer: B

The instance ID does not change with EC2 generation.

The change in hypervisor from Xen to Nitro may be the cause of the hardware ID change.
https://docs.aws.amazon.com/sap/latest/general/overview-sap-on-aws.html#:~:text=SAP%20Note%202113263,AWS%20Hardware%20ID

upvoted 2 times

**trashi** 2 years, 4 months ago

B is correct, see SAP note 1656250 - SAP on AWS: Support prerequisites

upvoted 2 times

**kk8s** 2 years, 4 months ago

**Selected Answer: A**

A

https://docs.aws.amazon.com/sap/latest/general/overview-sap-on-aws.html

upvoted 2 times

Question #21 *Topic 1*

A company's basis administrator is planning to deploy SAP on AWS in Linux. The basis administrator must set up the proper storage to store SAP HANA data and log volumes.
Which storage options should the basis administrator choose to meet these requirements? (Choose two.)

A. Amazon Elastic Block Store (Amazon EBS) Throughput Optimized HDD (st1)

B. Amazon Elastic Block Store (Amazon EBS) Provisioned OPS SSD (io1, io2)

C. Amazon S3

D. Amazon Elastic File System (Amazon EFS)

E. Amazon Elastic Block Store (Amazon EBS) General Purpose SSD (gp2, gp3)

**Suggested Answer:** *CD*

*Community vote distribution*


**[Removed]** 1 year, 11 months ago

**Selected Answer: BE**

https://docs.aws.amazon.com/sap/latest/sap-hana/hana-ops-storage-config.html

upvoted 2 times

**luckdragon** 1 year, 11 months ago

**Selected Answer: BE**

s3 & EFS are only suitable for backups, correct answers are B and E

upvoted 1 times

**easytoo** 2 years ago

b-e-b-e

upvoted 1 times

**schalke04** 2 years, 4 months ago

**Selected Answer: BE**

B and E

upvoted 3 times

**kk8s** 2 years, 5 months ago

**Selected Answer: BE**

BE as well

upvoted 3 times

**bigmuramura2** 2 years, 5 months ago

**Selected Answer: BE**

I think b,e

https://docs.aws.amazon.com/sap/latest/sap-hana/hana-ops-storage-config.html

upvoted 4 times

Question #22 *Topic 1*

A company has deployed a highly available SAP NetWeaver system on SAP HANA into a VPC. The system is distributed across multiple Availability Zones within a single AWS Region. SAP NetWeaver is running on SUSE Linux Enterprise Server for SAP. SUSE Linux Enterprise High Availability Extension is configured to protect SAP ASCS and ERS instances and uses the overlay IP address concept. The SAP shared files /sapmnt and /usr/sap/trans are hosted on an Amazon Elastic File System (Amazon EFS) file system.
The company needs a solution that uses already-existing private connectivity to the VPC. The SAP NetWeaver system must be accessible through the SAP GUI client tool.
Which solutions will meet these requirements? (Choose two.)

A. Deploy an Application Load Balancer. Configure the overlay IP address as a target.

B. Deploy a Network Load Balancer. Configure the overlay IP address as a target.

C. Use an Amazon Route 53 private zone. Create an A record that has the overlay IP address as a target.

D. Use AWS Transit Gateway. Configure the overlay IP address as a static route in the transit gateway route table. Specify the VPC as a target.

E. Use a NAT gateway. Configure the overlay IP address as a target.

**Suggested Answer:** *CE*

*Community vote distribution*

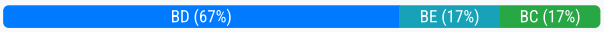


**ohcn** Highly Voted 2 years, 4 months ago

Selected Answer: BD

B and D

upvoted 5 times

**zzw890827** Most Recent 1 year, 10 months ago

Selected Answer: BC

AWS Transit Gateway is generally used for connecting multiple VPCs and on-premises networks. Adding a static route for the overlay IP does not specifically facilitate SAP GUI client access and would be more applicable for routing concerns between multiple networks.
so BC is right

upvoted 1 times

**easytoo** 2 years ago

b-d-b-d-b-d-b-d

upvoted 3 times

**easytoo** 2 years ago

By configuring the overlay IP address as a target in the NLB, incoming requests to the SAP NetWeaver system can be properly routed and load-balanced among the available ASCS and ERS instances.

upvoted 1 times

**[Removed]** 1 year, 11 months ago

B + D You can configure overlay IP routing with AWS Transit Gateway or Network Load Balancer
https://docs.aws.amazon.com/sap/latest/sap-hana/hana-ops-ha-dr.html

upvoted 2 times

**[Removed]** 1 year, 11 months ago

edit B,C

upvoted 1 times

**Hyperdanny** 2 years, 4 months ago

A + B are incorrect, since you wouldn't set the overlay IP adress as the target (OIP is the destination for network load balancers)
E seems incorrect: No need for the VPC instances to talk to the Internet.

So I am sticking with C + D

upvoted 3 times

**Hyperdanny** 2 years, 4 months ago

changing my mind: B+D,

upvoted 3 times

**SMALLAM** 2 years, 4 months ago

I think B&D

https://docs.aws.amazon.com/sap/latest/sap-hana/sap-oip-sap-on-aws-high-availability-setup.html

upvoted 4 times

**kk8s** 2 years, 5 months ago

**Selected Answer: BE**

B and E

NAT is need.

https://aws.amazon.com/blogs/awsforsap/vpc-subnet-zoning-patterns-for-sap-on-aws-part-3-internal-and-external-access/

upvoted 2 times

**forexamweb** 2 years, 5 months ago

**Selected Answer: BD**

I think B and D

upvoted 3 times

**bigmuramura2** 2 years, 5 months ago

**Selected Answer: BC**

i think b,c

upvoted 1 times

Question #23 *Topic 1*

A company is planning to move all its SAP applications to Amazon EC2 instances in a VPC. Recently, the company signed a multiyear contract with a payroll software-as-a-service (SaaS) provider. Integration with the payroll SaaS solution is available only through public web APIs. Corporate security guidelines state that all outbound traffic must be validated against an allow list. The payroll SaaS provider provides only fully qualified domain name (FQDN) addresses and no IP addresses or IP address ranges. Currently, an on-premises firewall appliance filters FQDNs. The company needs to connect an SAP Process Orchestration (SAP PO) system to the payroll SaaS provider.
What must the company do on AWS to meet these requirements?

A. Add an outbound rule to the security group of the SAP PO system to allow the FQDN of the payroll SaaS provider and deny all other outbound traffic.

B. Add an outbound rule to the network ACL of the subnet that contains the SAP PO system to allow the FQDN of the payroll SaaS provider and deny all other outbound traffic.

C. Add an AWS WAF web ACL to the VPAdd an outbound rule to allow the SAP PO system to connect to the FQDN of the payroll SaaS provider.

D. Add an AWS Network Firewall firewall to the VPC. Add an outbound rule to allow the SAP PO system to connect to the FQDN of the payroll SaaS provider.

**Suggested Answer:** *D*

*Community vote distribution*


**kaishin0527** 1 year, 11 months ago

Selected Answer: D

D: AWS Network Firewall is a managed service that makes it easy to deploy essential network protections for all of your Amazon Virtual Private Clouds (VPCs). The service can be configured to filter traffic based on fully qualified domain names (FQDN), which meets the requirement of the scenario.

upvoted 3 times

**[Removed]** 1 year, 11 months ago

Selected Answer: D

https://docs.aws.amazon.com/network-firewall/latest/developerguide/stateful-rule-groups-domain-names.html

upvoted 1 times

**easytoo** 2 years ago

d-d-d-d-dd--d-d-d-d

upvoted 1 times

**SMALLAM** 2 years, 4 months ago

I think D

upvoted 4 times

**Grillppl** 2 years, 4 months ago

Waf only incoming traffic and no fqdn filtering.
Nw fw inbound and outbound filtering and fqdn filtering supported. So it should be D

upvoted 3 times

**Balki** 2 years, 4 months ago

Selected Answer: D

FQDN filtering can be achieved only through Firewall https://aws.amazon.com/blogs/security/use-aws-network-firewall-to-filter-outbound-https-traffic-from-applications-hosted-on-amazon-eks/

upvoted 3 times

**forexamweb** 2 years, 4 months ago

Selected Answer: D

D maybe
https://aws.amazon.com/network-firewall/features#Web_filtering

upvoted 3 times

**kk8s** 2 years, 4 months ago

C for me

upvoted 1 times

**SONALID** 1 year, 8 months ago

web ACL rule cannot have FQDN

upvoted 1 times

Question #24 *Topic 1*

A company is planning to migrate its on-premises SAP application to AWS. The application runs on VMware vSphere. The SAP ERP Central Component (SAP ECC) server runs on an IBM Db2 database that is 2 TB in size. The company wants to migrate the database to SAP HANA. Which migration strategy will meet these requirements?

A. Use AWS Application Migration Service (CloudEndure Migration).

B. Use SAP Software Update Manager (SUM) Database Migration Option (DMO) with System Move.

C. Use AWS Server Migration Service (AWS SMS).

D. Use AWS Database Migration Service (AWS DMS).

**Suggested Answer:** *A*

*Community vote distribution*

B (93%) 7%

**kaishin0527** Highly Voted 1 year, 11 months ago

Selected Answer: B

B: The Database Migration Option (DMO) within the SAP Software Update Manager (SUM) is the recommended tool to convert your existing SAP ERP database to SAP HANA. This is because DMO is able to update the existing SAP system and migrate the database to SAP HANA in one step, hence reducing the total downtime required for the migration.

Option A: AWS Application Migration Service (CloudEndure Migration) doesn't support database transformation (such as converting IBM DB2 to SAP HANA).

Option C: AWS Server Migration Service (AWS SMS) is used for migrating on-premise servers to AWS, but it does not handle database conversions from IBM Db2 to SAP HANA.

Option D: AWS Database Migration Service (AWS DMS) helps you migrate databases to AWS easily and securely, but as of my knowledge cutoff in September 2021, it doesn't support SAP HANA as a target for migration.

upvoted 5 times

**[Removed]** Most Recent 1 year, 11 months ago

Selected Answer: C

VMWare classic use case for using AWS Server Migration Service (AWS SMS)

upvoted 1 times

**[Removed]** 1 year, 11 months ago

Edit B

upvoted 1 times

**easytoo** 2 years ago

b=b=b=b=b=b

upvoted 1 times

**schalke04** 2 years, 4 months ago

Selected Answer: B

B totally

upvoted 3 times

**kk8s** 2 years, 5 months ago

Selected Answer: B

B for Me.

https://docs.aws.amazon.com/sap/latest/sap-hana/migrating-hana-tools.html

upvoted 4 times

**bigmuramura2** 2 years, 5 months ago

Selected Answer: B

i think b

upvoted 2 times

Question #25 *Topic 1*

A company hosts multiple SAP applications on Amazon EC2 instances in a VPC. While monitoring the environment, the company notices that multiple port scans are attempting to connect to SAP portals inside the VPC. These port scans are originating from the same IP address block. The company must deny access to the VPC from all the offending IP addresses for the next 24 hours.
Which solution will meet this requirement?

A. Modify network ACLs that are associated with all public subnets in the VPC to deny access from the IP address block.

B. Add a rule in the security group of the EC2 instances to deny access from the IP address block.

C. Create a policy in AWS Identity and Access Management (IAM) to deny access from the IP address block.

D. Configure the firewall in the operating system of the EC2 instances to deny access from the IP address block.

**Suggested Answer:** *C*

*Community vote distribution*

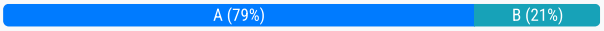


**kaishin0527** 1 year, 11 months ago

**Selected Answer: A**

A: Network Access Control Lists (ACLs) in Amazon VPC provide a layer of security for your VPC that act as a firewall for controlling traffic in and out of one or more subnets. You can add rules to your Network ACL to deny the traffic from specific IP address block.

upvoted 2 times

**[Removed]** 1 year, 11 months ago

**Selected Answer: B**

you only want to block the affected instance not all.

upvoted 2 times

**[Removed]** 1 year, 11 months ago

Edit A

upvoted 1 times

**easytoo** 2 years ago

a-a-a-a-a-a

Network ACLs (NACLs) act as a stateless firewall at the subnet level and control inbound and outbound traffic. They provide granular control over traffic flow, allowing or denying access based on IP address rules.

By modifying the NACLs associated with the public subnets in the VPC, the company can add a rule to deny access from the offending IP address block. This will effectively block any incoming or outgoing traffic from those IP addresses for the specified period of time (24 hours).

upvoted 2 times

**ohcn** 2 years, 4 months ago

**Selected Answer: A**

A I think

upvoted 3 times

**SMALLAM** 2 years, 4 months ago

A For me

upvoted 2 times

**schalke04** 2 years, 4 months ago

**Selected Answer: A**

A. Modify network ACLs that are associated with all public subnets in the VPC to deny access from the IP address block.

upvoted 2 times

**forexamweb** 2 years, 5 months ago

**Selected Answer: A**

A

Option B is incorrect. Security group does not support DENY rules.

upvoted 4 times

**kk8s** 2 years, 5 months ago

**Selected Answer: B**

B for me

upvoted 1 times

**kk8s** 2 years, 4 months ago

Sorry change to A

upvoted 2 times

Question #26 *Topic 1*

A company has deployed SAP workloads on AWS. The AWS Data Provider for SAP is installed on the Amazon EC2 instance where the SAP application is running. An SAP solutions architect has attached an IAM role to the EC2 instance with the following policy:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "AWSDataProvider1",
            "Effect": "Allow",
            "Action": [
                "EC2:DescribeInstances",
                "EC2:DescribeVolumes"
            ],
            "Resource": "*"
        },
        {
            "Sid": "AWSDataProvider2",
            "Effect": "Allow",
            "Action": "s3:GetObject",
            "Resource": [
                "arn:aws:s3:::aws-sap-data-provider/config.properties"
            ]
        }
    ]
}
```

The AWS Data Provider for SAP is not returning any metrics to the SAP application.
Which change should the SAP solutions architect make to the IAM permissions to resolve this issue?

A. Add the cloudwatch:ListMetrics action to the policy statement with Sid AWSDataProvider1.

B. Add the cloudwatch:GetMetricStatistics action to the policy statement with Sid AWSDataProvider1.

C. Add the cloudwatch:GetMetricStream action to the policy statement with Sid AWSDataProvider1.

D. Add the cloudwatch:DescribeAlarmsForMetric action to the policy statement with Sid AWSDataProvider1.

**Suggested Answer:** *A*

*Community vote distribution*

B (100%)

**kk8s** Highly Voted 2 years, 5 months ago

Selected Answer: B

B

https://docs.aws.amazon.com/sap/latest/general/data-provider-troubleshooting.html

upvoted 11 times

**[Removed]** Most Recent 1 year, 11 months ago

Selected Answer: B

GetMetricStatistics

upvoted 1 times

Question #27 *Topic 1*

A company wants to deploy an SAP HANA database on AWS by using AWS Launch Wizard for SAP. An SAP solutions architect needs to run a custom post-deployment script on the Amazon EC2 instance that Launch Wizard provisions.
Which actions can the SAP solutions architect take to provide the post-deployment script in the Launch Wizard console? (Choose two.)

A. Provide the FTP URL of the script.

B. Provide the HTTPS URL of the script on a web server.

C. Provide the Amazon S3 URL of the script.

D. Write the script inline.

E. Upload the script.

**Suggested Answer:** *BC*

*Community vote distribution*

CE (100%)

**kk8s** Highly Voted 2 years, 5 months ago

Selected Answer: CE

S3 and upload

https://catalog.us-east-1.prod.workshops.aws/workshops/754ba343-2704-404a-8abe-be7b21c4d9d5/en-US/800-other/802-prepostscript

upvoted 9 times

**[Removed]** Most Recent 1 year, 11 months ago

Selected Answer: CE

Upload --> S3

upvoted 1 times

**easytoo** 2 years ago

C. Providing the Amazon S3 URL of the script allows the SAP solutions architect to reference the location of the script stored in an Amazon S3 bucket. The Launch Wizard can then retrieve the script from the specified S3 URL during the deployment process.

E. Uploading the script directly in the Launch Wizard console provides the option to upload the script file directly from the local system. This allows the script to be included as part of the deployment package and be available for execution during the provisioning of the EC2 instance.

upvoted 3 times

**vivekmanasa** 2 years, 4 months ago

Selected Answer: CE

C and E

upvoted 2 times

Question #28 *Topic 1*

A company is planning to move its on-premises SAP HANA database to AWS. The company needs to migrate this environment to AWS as quickly as possible. An SAP solutions architect will use AWS Launch Wizard for SAP to deploy this SAP HANA workload.
Which combination of steps should the SAP solutions architect follow to start the deployment of this workload on AWS? (Choose three.)

A. Download the SAP HANA software.

B. Download the AWS CloudFormation template for the SAP HANA deployment.

C. Download and extract the SAP HANA software. Upload the SAP HANA software to an FTP server that Launch Wizard can access.

D. Upload the unextracted SAP HANA software to an Amazon S3 destination bucket. Follow the S3 file path syntax for the software in accordance with Launch Wizard recommendations.

E. Bring the operating system AMI by using the Bring Your Own Image (BYOI) model, or purchase the subscription for the operating system AMI from AWS Marketplace.

F. Create the SAP file system by using Amazon Elastic Block Store (Amazon EBS) before the deployment.

**Suggested Answer:** *BEF*

*Community vote distribution*

ADE (85%) Other

**forexamweb** Highly Voted 2 years, 5 months ago

Selected Answer: ADE

ADE

https://docs.aws.amazon.com/launchwizard/latest/userguide/launch-wizard-sap-setting-up.html
https://docs.aws.amazon.com/launchwizard/latest/userguide/launch-wizard-sap-structure.html

upvoted 8 times

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

Selected Answer: ADE

voting A,D,E

upvoted 1 times

**kaishin0527** 1 year, 11 months ago

Selected Answer: ADE

A,D,E: The first step is to download the SAP HANA software. This is a prerequisite for setting up an SAP HANA system.

Next, you need to upload the unextracted SAP HANA software to an Amazon S3 bucket. AWS Launch Wizard will use the software from this S3 bucket for the deployment.

Lastly, you need an AMI for the operating system. You can either use the BYOI model or purchase a subscription for the operating system AMI from AWS Marketplace.

AWS Launch Wizard does not require a pre-existing CloudFormation template (option B), nor does it need you to upload the SAP HANA software to an FTP server (option C). Launch Wizard creates necessary AWS resources, like the SAP file system, as part of the deployment process, so you don't need to create the SAP file system using Amazon EBS before the deployment (option F).

upvoted 3 times

**[Removed]** 1 year, 11 months ago

Selected Answer: ADE

voting A,D,E

upvoted 1 times

**JavaAWSLopez** 2 years ago

the company already have a HANA DB on premise, why using launch wizard? it will build a NEW Hana deployment, what about data from HANA on prem? no backup/restore?

upvoted 1 times

**ohcn** 2 years, 4 months ago

**Selected Answer: ADE**

ADE I think

upvoted 2 times

**SMALLAM** 2 years, 4 months ago

ADE for me

upvoted 2 times

**everydaysmile** 2 years, 4 months ago

**Selected Answer: ADE**

my answer

upvoted 2 times

**Balki** 2 years, 4 months ago

**Selected Answer: ABD**

cloudformation template is needed

upvoted 2 times

**kk8s** 2 years, 5 months ago

**Selected Answer: AE**

AE.

C wrong, not FTP

D wrong, no need to extract

https://docs.aws.amazon.com/launchwizard/latest/userguide/launch-wizard-sap-setting-up.html

upvoted 1 times

Question #29 *Topic 1*

A company wants to implement SAP HANA on AWS with the Multi-AZ deployment option by using AWS Launch Wizard for SAP. The solution will use SUSE Linux Enterprise High Availability Extension for the high availability deployment. An SAP solutions architect must ensure that all the prerequisites are met. The SAP solutions architect also must ensure that the user inputs to start the guided deployment of Launch Wizard are valid.
Which combination of steps should the SAP solutions architect take to meet these requirements? (Choose two.)

A. Before starting the Launch Wizard deployment, create the underlying Amazon Elastic Block Store (Amazon EBS) volume types to use for SAP HANA data and log volumes based on the performance requirements.

B. Use a value for the PaceMakerTag parameter that is not used by any other Amazon EC2 instances in the AWS Region where the system is being deployed.

C. Ensure that the virtual hostname for the SAP HANA database that is used for the SUSE Linux Enterprise High Availability Extension configuration is not used in any other deployed accounts.

D. Ensure that the VirtualIPAddress parameter is outside the VPC CIDR and is not being used in the route table that is associated with the subnets where primary and secondary SAP HANA instances will be deployed.

E. Before starting the Launch Wizard deployment, set up the SUSE Linux Enterprise High Availability Extension network configuration and security group.

**Suggested Answer:** *BE*

*Community vote distribution*

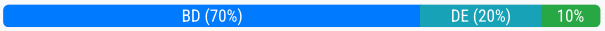


**[Removed]** 1 year, 11 months ago

**Selected Answer: BC**

Voting B,C

upvoted 1 times

**vivekmanasa** 2 years, 4 months ago

**Selected Answer: BD**

B and D

upvoted 3 times

**SMALLAM** 2 years, 4 months ago

I think BD

upvoted 2 times

**ohcn** 2 years, 4 months ago

B and D https://docs.aws.amazon.com/launchwizard/latest/userguide/launch-wizard-sap-deploying.html#launch-wizard-hana

upvoted 3 times

**forexamweb** 2 years, 4 months ago

**Selected Answer: BD**

I think BD.
In the document "SAP HANA on the AWS Cloud Quick Start Reference Deployment" (ARCHIVED) at the below URL, page 25, under "Requirements for Multi-AZ, Single-Node HA Scenarios," you will find the following

https://links.imagerelay.com/cdn/3404/ql/8327da608b7341f4ac2216c503116387/SAP-hana-on-AWS-cloud.pdf

####
SLES HAE and RHEL High Availability agents require that the Pacemaker tag and the overlay IP address you provide by setting deployment parameters can be uniquely identified. Therefore, you need to ensure the following:

– The value you provide for the PaceMakerTag parameter isn't being used by any other EC2 instances in your account, in the AWS Region where you are deploying the Quick Start.
– The IP address you provide for the VirtualIPAddress parameter is outside the VPC CIDR and isn't being used in the route table associated with the

subnets where primary and secondary HANA instances will be deployed.
####
upvoted 4 times

**forexamweb** 2 years, 4 months ago
Furthermore, the latest document (AWS Launch Wizard User Guide) contains a similar statement.

1. Go to https://docs.aws.amazon.com/launchwizard/latest/userguide/launch-wizard-sap-deploying.html#launch-wizard-hana
2. Select tab "High availability deployment"

You can see the following statements.

####
Overlay IP address. Enter the overlay IP address to assign to the active node. The IP address should be outside of the VPC CIDR and must not be used by any other HA cluster. It is configured to always point to the active SAP HANA node.

Pacemaker tag name. Enter the tag to assign to each EC2 instance. This tag is used by the pacemaker component of SLES HAE and RHEL for SAP high availability solutions and must not be used by any other EC2 instance in your account.
####
upvoted 4 times

**forexamweb** 2 years, 4 months ago
So, BD
upvoted 3 times

**schalke04** 2 years, 4 months ago
Selected Answer: DE
D and E
upvoted 1 times

**Balki** 2 years, 4 months ago
Selected Answer: DE
Question & Answers are vague. D&E are closest
upvoted 1 times

**Hyperdanny** 2 years, 4 months ago
My answer is D + E
upvoted 1 times

Question #30 *Topic 1*

A company that has SAP workloads on premises plans to migrate an SAP environment to AWS. The company is new to AWS and has no prior setup. The company has the following requirements:
The application server and database server must be placed in isolated network configurations.
SAP systems must be accessible to the on-premises end users over the internet.
The cost of communications between the application server and the database server must be minimized.
Which combination of steps should an SAP solutions architect take to meet these requirements? (Choose two.)

A. Configure a Network Load Balancer for incoming connections from end users.

B. Set up an AWS Site-to-Site VPN connection between the company's on-premises network and AWS.

C. Separate the application server and the database server by using different VPCs.

D. Separate the application server and the database server by using different subnets and network security groups within the same VPC.

E. Set up an AWS Direct Connect connection with a private VIF between the company's on-premises network and AWS.

**Suggested Answer:** *BE*

*Community vote distribution*

BD (89%) 11%

**MKM** Highly Voted 2 years, 4 months ago
BD is correct!
upvoted 10 times

**Technolord** Most Recent 1 year, 10 months ago
I wonder who decide the answers here? How can it be B and E?? What is the poin of setting up both a DC and a site-to-site VPN?
upvoted 1 times

**[Removed]** 1 year, 11 months ago
Selected Answer: BD
I think B,D
upvoted 1 times

**easytoo** 2 years ago
b-d-b-d-b-d-b-d
upvoted 1 times

**Worldclassarchitect** 2 years, 4 months ago
Shouldn't it be C instead of D.

D says 'Network Security Groups' - there is no such thing. It is either Network Access Control List at the Subnet Level or Security groups at the instance levels.
upvoted 1 times

**SMALLAM** 2 years, 4 months ago
BD for sure
upvoted 2 times

**asharma45** 2 years, 4 months ago
AD since connection is over internet
upvoted 2 times

**schalke04** 2 years, 4 months ago
Selected Answer: BD
Direct connect is expensive. NLB is not relevant.
VPN is the cheapest option for the new customer.
upvoted 3 times

**matakuyy2** 2 years, 4 months ago
Selected Answer: AD

From on-premises, AWS Site-to-Site VPN is not required since it is via the Internet.
upvoted 1 times

**Balki** 2 years, 4 months ago

NLB is the costliest. If we need to minimize cost, we should go with BD
upvoted 3 times

**kk8s** 2 years, 5 months ago

**Selected Answer: BD**

BD for me
upvoted 4 times

Question #31 *Topic 1*

A company is running its SAP workload on AWS. The company's security team has implemented the following requirements:
All Amazon EC2 instances for SAP must be SAP certified instance types.
Encryption must be enabled for all Amazon S3 buckets and Amazon Elastic Block Store (Amazon EBS) volumes.
AWS CloudTrail must be activated.
SAP system parameters must be compliant with business rules.
Detailed monitoring must be enabled for all instances.
The company wants to develop an automated process to review the systems for compliance with the security team's requirements. The process also must provide notification about any deviation from these standards.
Which solution will meet these requirements?

A. Use AWS AppConfig to model configuration data in an AWS Systems Manager Automation runbook. Schedule this Systems Manager Automation runbook to monitor for compliance with all the requirements. Integrate AWS AppConfig with Amazon CloudWatch for notification purposes.

B. Use AWS Config managed rules to monitor for compliance with all the requirements. Use Amazon EventBridge (Amazon CloudWatch Events) and Amazon Simple Notification Service (Amazon SNS) for email notification when a resource is flagged as noncompliant.

C. Use AWS Trusted Advisor to monitor for compliance with all the requirements. Use Trusted Advisor preferences for email notification when a resource is flagged as noncompliant.

D. Use AWS Config managed rules to monitor for compliance with the requirements, except for the SAP system parameters. Create AWS Config custom rules to validate the SAP system parameters. Use Amazon EventBridge (Amazon CloudWatch Events) and Amazon Simple Notification Service (Amazon SNS) for email notification when a resource is flagged as noncompliant.

**Suggested Answer:** *D*

*Community vote distribution*

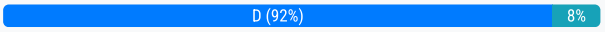


**kaishin0527** 1 year, 11 months ago

Selected Answer: D

D: AWS Config managed rules can monitor AWS resources for compliance with specified configurations. However, AWS Config does not have built-in functionality to monitor SAP system parameters, so you would need to create custom rules for this purpose. AWS Config can then use Amazon EventBridge (formerly known as Amazon CloudWatch Events) to trigger notifications via Amazon SNS when a resource is found to be noncompliant. This solution provides the needed automation and compliance review capabilities.

upvoted 1 times

**[Removed]** 1 year, 11 months ago

Selected Answer: D

Obvious D,

upvoted 1 times

**easytoo** 2 years ago

d-d-d-d-d-d

upvoted 2 times

**schalke04** 2 years, 4 months ago

Selected Answer: D

D looks good

upvoted 4 times

**Kiran1982** 2 years, 4 months ago

Selected Answer: D

https://aws.amazon.com/blogs/awsforsap/audit-your-sap-systems-with-aws-config-part-ii/

upvoted 4 times

**forexamweb** 2 years, 5 months ago

Selected Answer: D

D

https://aws.amazon.com/blogs/awsforsap/audit-your-sap-systems-with-aws-config-part-i/

https://aws.amazon.com/blogs/awsforsap/audit-your-sap-systems-with-aws-config-part-ii/
upvoted 1 times

**kk8s** 2 years, 5 months ago

**Selected Answer: B**

B for me.

https://aws.amazon.com/blogs/awsforsap/audit-your-sap-systems-with-aws-config-part-i/
upvoted 1 times

Question #32 *Topic 1*

A company is hosting its SAP workloads on AWS. An SAP solutions architect is designing high availability architecture for the company's production SAP S/4HANA and SAP BW/4HANA workloads. These workloads have the following requirements:
Redundant SAP application servers that consist of a primary application server (PAS) and an additional application server (AAS)
ASCS and ERS instances that use a failover cluster
Database high availability with a primary DB instance and a secondary DB instance
How should the SAP solutions architect design the architecture to meet these requirements?

A. Deploy ASCS and ERS cluster nodes in different subnets within the same Availability Zone. Deploy the PAS instance and AAS instance in different subnets within the same Availability Zone. Deploy the primary DB instance and secondary DB instance in different subnets within the same Availability Zone. Deploy all the components in the same VPC.

B. Deploy ASCS and ERS cluster nodes in different subnets within the same Availability Zone. Deploy the PAS instance and AAS instance in different subnets within the same Availability Zone. Deploy the primary DB instance and secondary DB instance in different subnets within the same Availability Zone. Deploy the ASCS instance, PAS instance, and primary DB instance in one VPC. Deploy the ERS instance, AAS instance, and secondary DB instance in a different VPC.

C. Deploy ASCS and ERS cluster nodes in different subnets across two Availability Zones. Deploy the PAS instance and AAS instance in different subnets across two Availability Zones. Deploy the primary DB instance and secondary DB instance in different subnets across two Availability Zones. Deploy all the components in the same VPC.

D. Deploy ASCS and ERS cluster nodes in different subnets across two Availability Zones. Deploy the PAS instance and AAS instance in different subnets across two Availability Zones. Deploy the primary DB instance and secondary DB instance in different subnets across two Availability Zones. Deploy the ASCS instance, PAS instance, and primary DB instance in one VPC. Deploy the ERS instance, AAS instance, and secondary DB instance in a different VPC.

**Suggested Answer:** *D*

*Community vote distribution*


**muzavor** 1 year, 6 months ago

C

C looks right . We donot want to have two VPC within the same region . so this makes sense

upvoted 1 times

**kaishin0527** 1 year, 11 months ago

**Selected Answer: C**

C: This setup meets all the high availability requirements of the SAP workloads. The ASCS and ERS instances, which are part of the central services layer of SAP system, are deployed in a failover cluster across different subnets and Availability Zones, ensuring high availability and failover protection. The PAS and AAS are also deployed across different subnets and Availability Zones, enabling redundancy and load distribution. The primary and secondary database instances are also deployed across different subnets and Availability Zones, providing high availability for the database layer. All components are deployed within the same VPC, which enables efficient network traffic routing between them.

upvoted 2 times

**[Removed]** 1 year, 11 months ago

**Selected Answer: C**

voting c

upvoted 1 times

**easytoo** 2 years ago

c-c-c-c-c-c

upvoted 1 times

**SMALLAM** 2 years, 4 months ago

C, D is incorrect because Deploying all the components in the same VPC makes it easier to manage network communication and security for the SAP workloads.

upvoted 2 times

**schalke04** 2 years, 4 months ago

**Selected Answer: C**

100% C

upvoted 3 times

**Hyperdanny** 2 years, 4 months ago

Voting C

upvoted 2 times

**MKM** 2 years, 4 months ago

The answer is C.

upvoted 2 times

**kk8s** 2 years, 5 months ago

**Selected Answer: C**

C for me.

https://www.youtube.com/watch?v=0VyvrE2fvmc

upvoted 4 times

Question #33 *Topic 1*

A company has deployed SAP HANA in the AWS Cloud. The company needs its SAP HANA database to be highly available. An SAP solutions architect has deployed the SAP HANA database in separate Availability Zones in a single AWS Region. SUSE Linux Enterprise High Availability Extension is configured with an overlay IP address. The overlay IP resource agent has the following IAM policy:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "oip1",
            "Effect": "Allow",
            "Action": "ec2:AssociateRouteTable",
            "Resource": "arn:aws:ec2:us-east-1:111111111111:route-table/rtb-XYZ"
        },
        {
            "Sid": "oip2",
            "Effect": "Allow",
            "Action": "ec2:DescribeRouteTables",
            "Resource": "*"
        }
    ]
}
```

During a test of failover, the SAP solutions architect finds that the overlay IP address does not change to the secondary Availability Zone. Which change should the SAP solutions architect make in the policy statement for Sid oip1 to fix this error?

A. Change the Action element to ec2:CreateRoute.

B. Change the Action element to ec2:ReplaceRoute.

C. Change the Action element to ec2:ReplaceRouteTableAssociation.

D. Change the Action element to ec2:ReplaceTransitGatewayRoute.

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**[Removed]** 1 year, 11 months ago

**Selected Answer: B**

ec2:ReplaceRoute

upvoted 2 times

**easytoo** 2 years ago

b-b-b-b-b-b

upvoted 1 times

**schalke04** 2 years, 4 months ago

**Selected Answer: B**

B. Change the Action element to ec2:ReplaceRoute.

upvoted 2 times

**trashi** 2 years, 4 months ago

B

https://docs.aws.amazon.com/sap/latest/sap-hana/sap-hana-on-aws-cluster-configuration-prerequisites.html

upvoted 3 times

**kk8s** 2 years, 5 months ago

**Selected Answer: B**

B for me.

https://docs.aws.amazon.com/sap/latest/sap-hana/sap-hana-on-aws-oip.html

upvoted 4 times

Question #34 *Topic 1*

A company wants to improve the RPO and RTO for its SAP disaster recovery (DR) solution by running the DR solution on AWS. The company is running SAP ERP Central Component (SAP ECC) on SAP HANA. The company has set an RPO of 15 minutes and an RTO of 4 hours.
The production SAP HANA database is running on a physical appliance that has x86 architecture. The appliance has 1 TB of memory, and the SAP HANA global allocation limit is set to 768 GB. The SAP application servers are running as VMs on VMware, and they store data on an NFS file system. The company does not want to change any existing SAP HANA parameters that are related to data and log backup for its on-premises systems.
What should an SAP solutions architect do to meet the DR objectives MOST cost-effectively?

A. For the SAP HANA database, change the log backup frequency to 5 minutes. Move the data and log backups to Amazon S3 by using the AWS CLI or AWS DataSync. Launch the SAP HANA database. For the SAP application servers, export the VMs as AMIs by using the VM Import/Export feature from AWS. For NFS file shares /sapmnt and /usr/sap/trans, establish real-time synchronization from DataSync to Amazon Elastic File System (Amazon EFS).

B. For the SAP HANA database, change the log backup frequency to 5 minutes. Move the data and log backups to Amazon S3 by using AWS Storage Gateway File Gateway. For the SAP application servers, export the VMs as AMIs by using the VM Import/Export feature from AWS. For NFS file shares /sapmnt and /usr/sap/trans, establish real-time synchronization from AWS DataSync to Amazon Elastic File System (Amazon EFS).

C. For the SAP HANA database, SAP application servers, and NFS file shares, use CloudEndure Disaster Recovery to replicate the data continuously from on premises to AWS. Use CloudEndure Disaster Recovery to launch target instances in the event of a disaster.

D. For the SAP HANA database, use a smaller SAP certified Amazon EC2 instance. Use SAP HANA system replication with ASYNC replication mode to replicate the data continuously from on premises to AWS. For the SAP application servers, use CloudEndure Disaster Recovery for continuous data replication. For NFS file shares /sapmnt and /usr/sap/trans, establish real-time synchronization from AWS DataSync to Amazon Elastic File System (Amazon EFS).

**Suggested Answer:** *A*

*Community vote distribution*

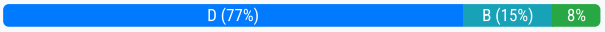


**schalke04** Highly Voted 2 years, 4 months ago

Selected Answer: D

ASYNC for HANA DB.
DRS for App

upvoted 6 times

**DigvijayGhosh** Most Recent 1 year, 7 months ago

D is correct, meets the requirement

upvoted 1 times

**Dhieraj** 1 year, 10 months ago

Selected Answer: D

D meets all the requirements. A and B gets eliminated as they are talking about parameter change which is against the ask in question. Cloud endure is block level replication so can not guarantee consistency of DB. Hence C gets eliminated.

upvoted 2 times

**kaishin0527** 1 year, 11 months ago

Selected Answer: D

D: The scenario mentions an RPO of 15 minutes and an RTO of 4 hours, so the company needs a disaster recovery solution that provides data replication to meet these objectives. The solution proposed in option D meets these requirements most cost-effectively. It includes SAP HANA system replication for the SAP HANA database, CloudEndure Disaster Recovery for the SAP application servers, and AWS DataSync for NFS file shares. This approach ensures the company can achieve the desired RPO and RTO without changing any existing SAP HANA parameters related to data and log backup for on-premises systems.

upvoted 1 times

**[Removed]** 1 year, 11 months ago

Selected Answer: C

VOTING C

upvoted 1 times

**[Removed]** 1 year, 11 months ago

changing to D

upvoted 1 times

**easytoo** 2 years ago

B. Changing the log backup frequency to 5 minutes ensures a lower RPO (Recovery Point Objective) by reducing the amount of potential data loss in the event of a disaster.

Moving the data and log backups to Amazon S3 using AWS Storage Gateway File Gateway allows for cost-effective storage of the backups in a durable and scalable manner.

Exporting the SAP application servers as Amazon Machine Images (AMIs) using the VM Import/Export feature enables their easy deployment in AWS in the event of a disaster, reducing the RTO (Recovery Time Objective).

Establishing real-time synchronization from on-premises NFS file shares (/sapmnt and /usr/sap/trans) to Amazon EFS using AWS DataSync ensures that the data remains up to date in the target environment.

upvoted 1 times

**SuiR** 2 years, 2 months ago

**Selected Answer: D**

it mentioned that "The company does not want to change any existing SAP HANA parameters that are related to data and log backup", A and B change the frequency of log backup, so not correct.
For C, it does not mentioned clearly, normally SAP HANA DB need an EC2 instance to be installed for data replication, it does not mentioned EC2 instance size, so assume it will use normal instance type which is not cost efficient.
So D is the better one , launch a smaller EC2 for hana DB and switch off preload.

upvoted 1 times

**blanco750** 2 years, 3 months ago

**Selected Answer: B**

Its between A or B. B looks right because in A it mentions creating DB which probabaly is not required in DR scenario.

upvoted 2 times

**CloudRover** 2 years, 3 months ago

i think its A because in B it uses both file gateway and datasync agent. This requires 2 agents to be set up which is not cost effective.

upvoted 1 times

**SuiR** 2 years, 2 months ago

it mentioned "The company does not want to change any existing SAP HANA parameters that are related to data and log backup ", so A and B should be not correct, I will choose D.

upvoted 1 times

Question #35 *Topic 1*

A company is planning to migrate its on-premises SAP applications to AWS. The applications are based on Windows operating systems. A file share stores the transport directories and third-party application data on the network-attached storage of the company's on-premises data center. The company's plan is to lift and shift the SAP applications and the file share to AWS. The company must follow AWS best practices for the migration.
Which AWS service should the company use to host the transport directories and third-party application data on AWS?

A. Amazon Elastic Block Store (Amazon EBS)

B. AWS Storage Gateway

C. Amazon Elastic File System (Amazon EFS)

D. Amazon FSx for Windows File Server

**Suggested Answer:** *C*

*Community vote distribution*

D (100%)

**G4Exams** 1 year, 9 months ago

Selected Answer: D

They are talking about Windows Server so it is D. FSx is the AWS service to go for.

upvoted 3 times

**kaishin0527** 1 year, 11 months ago

Selected Answer: D

D: Amazon FSx for Windows File Server is a fully managed service that provides cost-effective, highly reliable, and scalable file storage that is accessible over the industry-standard Server Message Block (SMB) protocol. It is built on Windows Server and offers a rich set of enterprise storage capabilities with the scalability, reliability, and low cost of AWS. Amazon FSx integrates with AWS managed Active Directory, allowing you to access your file systems using your existing Windows-based environments. It's an ideal choice for use cases like lift-and-shift enterprise applications, home directories, and software development.

upvoted 3 times

**[Removed]** 1 year, 11 months ago

Selected Answer: D

Keyword Windows = Amazon FSx for Windows File Server

upvoted 2 times

**easytoo** 2 years ago

d-d-d-d-d-d

upvoted 1 times

**blanco750** 2 years, 3 months ago

Selected Answer: D

FSX is for Windows. D is correct

upvoted 1 times

**Balki** 2 years, 4 months ago

Selected Answer: D

https://aws.amazon.com/blogs/awsforsap/how-to-setup-sap-netweaver-on-windows-mscs-for-sap-ascs-ers-on-aws-using-amazon-fsx/

upvoted 3 times

**kk8s** 2 years, 5 months ago

Selected Answer: D

D i think

upvoted 2 times

Question #36 *Topic 1*

A company hosts an SAP HANA database on an Amazon EC2 instance in the us-east-1 Region. The company needs to implement a disaster recovery (DR) site in the us-west-1 Region. The company needs a cost-optimized solution that offers a guaranteed capacity reservation, an RPO of less than 30 minutes, and an RTO of less than 30 minutes.
Which solution will meet these requirements?

A. Deploy a single EC2 instance to support the secondary database in us-west-1 with additional storage. Use this secondary database instance to support QA and production. Configure the primary SAP HANA database in us-east-1 to constantly replicate the data to the secondary SAP HANA database in us-west-1 by using SAP HANA system replication with preload off. During DR, shut down the QA SAP HANA instance and restart the production services at the secondary site.

B. Deploy a secondary staging server on an EC2 instance in us-west-1. Use CloudEndure Disaster Recovery to replicate changes at the database level from us-east-1 to the secondary staging server on an ongoing basis. During DR, initiate cutover, increase the size of the secondary EC2 instance to match the primary EC2 instance, and start the secondary EC2 instance.

C. Set up the primary SAP HANA database in us-east-1 to constantly replicate the data to a secondary SAP HANA database in us-west-1 by using SAP HANA system replication with preload on. Keep the secondary SAP HANA instance as a hot standby that is ready to take over in case of failure.

D. Create an SAP HANA database AMI by using Amazon Elastic Block Store (Amazon EBS) snapshots. Replicate the database and log backup files from a primary Amazon S3 bucket in us-east-1 to a secondary S3 bucket in us-west-1. During DR, launch the EC2 instance in us-west-1 based on AMIs that are replicated. Update host information. Download database and log backups from the secondary S3 bucket. Perform a point-in-time recovery.

**Suggested Answer:** *C*

*Community vote distribution*


**schalke04** Highly Voted 2 years, 4 months ago

Selected Answer: A

apply QA on DR site will satisfy cost-optimized solution.

upvoted 5 times

**blanco750** Highly Voted 2 years, 3 months ago

Selected Answer: B

Its B !. CloudEndure Disaster Recovery continuously replicates your machines (including operating system, system state configuration, databases, applications, and files) into a low-cost staging area in your target AWS account and preferred Region.
https://aws.amazon.com/blogs/awsforsap/sap-disaster-recovery-solution-using-cloudendure-part-1-failover/

upvoted 5 times

**acethetest1000** Most Recent 1 year, 5 months ago

Selected Answer: B

Running a QA system in another region will implicate on NFS issues. Hence I think the best option is B.

upvoted 1 times

**DigvijayGhosh** 1 year, 7 months ago

Selected Answer: A

"A" looks correct

upvoted 1 times

**[Removed]** 1 year, 7 months ago

Selected Answer: A

A is the answer. It was a close call between A & B until you read the requirement for capacity reservation, which B only promises a smaller instance, and they could face capacity issues when the EC2 instance is scalled to the prod size one.

upvoted 2 times

**acethetest1000** 1 year, 8 months ago

Selected Answer: A

I think A is the only cost-effective option that offers guaranteed capacity reservation. CloudEndure would work but it doesn't solve the capacity reservation requirement as it uses a staging area and launch the instances in case of a disaster. What it there is no spare capacity in the DR region?

upvoted 2 times

**student2020** 1 year, 8 months ago

Only C provide guaranteed capacity reservation as required by the question. With option B, during DR you can try to launch instance an get an error that there is no capacity in the AZ or region of the desired instance type.

upvoted 2 times

**SONALID** 1 year, 8 months ago

You can use CloudEndure Disaster Recovery to protect your most critical databases, including Oracle, MySQL, and Microsoft SQL Server, as well as enterprise applications such as SAP.

For HANA DB, Async replication is the best option to achieve RTO and RPO of less than 30mins. So option A is correct.

upvoted 1 times

**zzw890827** 1 year, 10 months ago

**Selected Answer: B**

Using a single EC2 instance in us-west-1 for both QA and production may not be ideal because it doesn't offer guaranteed capacity reservation for DR, which is one of the requirements. Shutting down QA to enable production may also lead to other issues and is not optimal.

upvoted 3 times

**kaishin0527** 1 year, 11 months ago

**Selected Answer: A**

A: This approach provides a cost-effective solution by utilizing a single EC2 instance to serve dual purposes - supporting both QA and production. By using SAP HANA system replication with preload off, it ensures a constant replication of data from the primary database to the secondary one. In case of a disaster recovery scenario, the QA SAP HANA instance can be shut down and production services can be restarted at the secondary site. This meets the RPO and RTO requirements of less than 30 minutes.

upvoted 2 times

**[Removed]** 1 year, 11 months ago

**Selected Answer: C**

voting C

upvoted 2 times

**juanvepe** 2 years ago

CloudEndure Disaster Recovery continuously replicates your machines (including operating system, system state configuration, databases, applications, and files) into a low-cost staging area in your target AWS account and preferred Region. In the case of a disaster, you can instruct CloudEndure Disaster Recovery to automatically launch thousands of your machines in their fully provisioned state in minutes. By replicating your machines into a low-cost staging area while still being able to launch fully provisioned machines within minutes, CloudEndure Disaster Recovery can significantly reduce the cost of your disaster recovery infrastructure. The two key concepts when it comes to DR planning are Recovery Time Objective (RTO) and Recovery Point Objective (RPO). RTO is the maximum period you want your systems to be unavailable due to an outage. RPO refers to the point of data processing you wish to recover to if there is a disaster. The following diagram illustrates the correlation of RTO and RPO:

upvoted 2 times

**anttan** 2 years, 4 months ago

B is the answer

upvoted 4 times

**Grillppl** 2 years, 4 months ago

I think A

upvoted 4 times

Question #37 *Topic 1*

An SAP solutions architect is leading the SAP basis team for a company. The company's SAP landscape includes SAP HANA database instances for the following systems: sandbox, development, quality assurance test (QAT), system performance test (SPT), and production. The sandbox, development, and QAT systems are running on Amazon EC2 On-Demand Instances. The SPT and production systems are running on EC2 Reserved instances. All the EC2 instances are using Provisioned IOPS SSO (io2) Amazon Elastic Block Store (Amazon EBS) volumes.
The entire development team is in the same time zone and works from 8 AM to 6 PM. The sandbox system is for research and testing that are not critical. The SPT and production systems are business critical. The company runs load-testing jobs and stress-testing jobs on the QAT systems overnight to reduce testing duration. The company wants to optimize infrastructure cost for the existing AWS resources.
How can the SAP solutions architect meet these requirements with the LEAST amount of administrative effort?

A. Use a Spot Fleet instead of the Reserved Instances and On-Demand Instances.

B. Use Amazon EventBridge (Amazon CloudWatch Events) and Amazon CloudWatch alarms to stop the development and sandbox EC2 instances from 7 PM every night to 7 AM the next day.

C. Make the SAP basis team available 24 hours a day, 7 days a week to use the AWS CLI to stop and start the development and sandbox EC2 instances manually.

D. Change the EBS volume type to Throughput Optimized HDD (st1) for the /hana/data and /hana/log file systems for the production and non-production SAP HANA databases.

**Suggested Answer:** *D*

*Community vote distribution*

B (100%)

**kaishin0527** 1 year, 11 months ago

**Selected Answer: B**

B: Amazon EventBridge can automate your AWS services and respond to system events such as application availability issues or resource changes. Events from AWS services are delivered to EventBridge in near-real time. You can write simple rules to indicate which events are of interest to you and what automated actions to take when an event matches a rule. Amazon CloudWatch Alarms watches a single metric over a time period you specify and performs one or more actions based on the value of the metric relative to a threshold over time.

Therefore, in this scenario, you can automate stopping of the development and sandbox EC2 instances during the off-peak hours to save costs without requiring the SAP basis team to manually stop and start the instances, which will save time and reduce administrative effort.

upvoted 4 times

**[Removed]** 1 year, 11 months ago

**Selected Answer: B**

voting - logical option

upvoted 2 times

**easytoo** 2 years ago

b-b-b-b-b-b

upvoted 3 times

**blanco750** 2 years, 3 months ago

**Selected Answer: B**

B is correct

upvoted 1 times

**schalke04** 2 years, 4 months ago

**Selected Answer: B**

B is correct

upvoted 3 times

**Hyperdanny** 2 years, 4 months ago

I am voting for B.
Using HDD for /data in a Prod environment doesn't sound look like a good idea (Option D).

upvoted 2 times

Question #38 *Topic 1*

A company is hosting an SAP HANA database on AWS. The company is automating operational tasks, including backup and system refreshes. The company wants to use SAP HANA Studio to perform data backup of an SAP HANA tenant database to a backint interface. The SAP HANA database is running in multi-tenant database container (MDC) mode. The company receives the following error message during an attempt to perform the backup:

```
Could not start backup for system <SID> DBC: [447]: backup could not be completed: [110091] Invalid path selection
for data backup using backint: usr/sap/<SID>/SYS/global/hdb/backint/COMPLETE_DATA_BACKUP must start with
/usz/sap/<SID>/SYS/global/hdb/backint/DB_<TENANT>.
```

What should an SAP solutions architect do to resolve this issue?

A. Set the execute permission for AWS Backint agent binary aws-backint-agent and for the launcher script aws-backint-agent-launcher.sh in the installation directory.

B. Verify the installation steps. Create symbolic links (symlinks).

C. Ensure that the catalog_backup_using_backint SAP HANA parameter is set to true. Ensure that the data_backup_parameter_file and log_backup_parameter_file parameters have the correct path location in the global.ini file.

D. Add the SAP HANA system to SAP HANA Studio. Select multiple container mode, and then try to initiate the backup again.

**Suggested Answer:** *D*

*Community vote distribution*

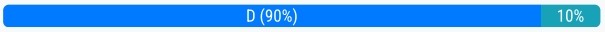


**awsmonster** 1 year, 4 months ago

Selected Answer: D

Agree with Nots, thanks for the link

https://docs.aws.amazon.com/ja_jp/sap/latest/sap-hana/aws-backint-agent-troubleshooting.html

upvoted 2 times

**r0se_mary** 1 year, 9 months ago

https://docs.aws.amazon.com/sap/latest/sap-hana/aws-backint-agent-troubleshooting.html

upvoted 2 times

**[Removed]** 1 year, 11 months ago

Selected Answer: B

voting B

upvoted 1 times

**[Removed]** 1 year, 11 months ago

Edit D

upvoted 1 times

**easytoo** 2 years ago

d-d-d-d-dd-

upvoted 1 times

**blanco750** 2 years, 3 months ago

Selected Answer: D

D is correct !

upvoted 1 times

**schalke04** 2 years, 4 months ago

Selected Answer: D

D:

https://me.sap.com/notes/0002512397

upvoted 2 times

**Nots** 2 years, 4 months ago

Selected Answer: D

I agree with D.

See "Problem: The following error is displayed when initiating a backup from the SAP HANA console:" section following URL.

https://docs.aws.amazon.com/ja_jp/sap/latest/sap-hana/aws-backint-agent-troubleshooting.html
upvoted 4 times

Question #39 *Topic 1*

A company is planning to migrate its on-premises SAP ERP Central Component (SAP ECC) system on SAP HANA to AWS. Each month, the system experiences two peaks in usage. The first peak is on the 21st day of the month when the company runs payroll. The second peak is on the last day of the month when the company processes and exports credit data. Both peak workloads are of high importance and cannot be rescheduled.
The current SAP ECC system has six application servers, all of a similar size. During normal operation outside of peak usage, four application servers would suffice.
Which purchasing option will meet the company's requirements MOST cost-effectively on AWS?

A. Four Reserved Instances and two Spot Instances

B. Six On-Demand Instances

C. Six Reserved Instances

D. Four Reserved Instances and two On-Demand Instances

**Suggested Answer:** *C*

*Community vote distribution*

D (100%)

**schalke04** Highly Voted 2 years, 4 months ago

Selected Answer: D

d:
A. Four Reserved Instances and two on demand Instances would meet the company's requirements most cost-effectively on AWS. The company can use four Reserved Instances for the normal operation and two on demand Instances for the peak workloads during the 21st day of the month and the last day of the month. This will provide the required capacity at the lowest cost.

upvoted 6 times

**DigvijayGhosh** Most Recent 1 year, 7 months ago

Selected Answer: D

D is correct

upvoted 1 times

**kaishin0527** 1 year, 11 months ago

Selected Answer: D

D: Reserved Instances are suitable for predictable workloads and provide a significant discount compared to On-Demand instance pricing. In this case, because four servers are consistently needed, it's cost-effective to use Reserved Instances for these.

However, because the other two servers are only needed during peak usage, it's more cost-effective to use On-Demand Instances, which let you pay for compute capacity by the hour or second (minimum of 60 seconds), with no long-term commitments.

upvoted 4 times

**[Removed]** 1 year, 11 months ago

Selected Answer: D

option D

upvoted 2 times

**easytoo** 2 years ago

d-d-d-d-d-d

upvoted 1 times

**easytoo** 2 years ago

no it's a-a-a-a-a-a-a-a-a-

upvoted 1 times

**easytoo** 2 years ago

jajaja no it's d-d-d-d-d-d

upvoted 1 times

**Hyperdanny** 2 years, 4 months ago

Why not D? Using 4 reserved instances and 2 additional On Demand when needed.

upvoted 3 times

Question #40 *Topic 1*

A company has an SAP environment that runs on AWS. The company wants to enhance security by restricting Amazon EC2 Instance Metadata Service (IMDS) to IMDSv2 only. The company's current configuration option supports both IMDSv1 and IMDSv2. The security enhancement must not create an SAP outage.
What should the company do before it applies the security enhancement on EC2 instances that are running the SAP environment?

A. Ensure that the SAP kernel versions are 7.45 or later.

B. Ensure that the EC2 instances are Nitro based.

C. Ensure that the AWS Data Provider for SAP is installed on each EC2 instance.

D. Stop the EC2 instances.

**Suggested Answer:** *A*

*Community vote distribution*

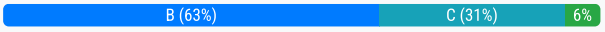


**anttan** Highly Voted 2 years, 4 months ago
Answer is A.

Before applying the security enhancement on EC2 instances running the SAP environment to restrict IMDS to IMDSv2 only, the company should ensure that the SAP kernel versions are 7.45 or later. This is because kernel version 7.45 and later supports the IMDSv2 protocol, while earlier versions only support IMDSv1. If the company applies the security enhancement before upgrading to kernel version 7.45 or later, it could result in an SAP outage or other issues.
upvoted 6 times

**awsmonster** Most Recent 1 year, 4 months ago
Selected Answer: C
Agreed with geoakes !
upvoted 1 times

**geoakes** 1 year, 5 months ago
Selected Answer: C
Why C: https://aws.amazon.com/about-aws/whats-new/2021/05/aws-data-provider-sap-version-4-0-now-available/

Not A or B, since this is a 'SAP environment that runs on AWS' already
Not D as that is downtime
upvoted 2 times

**kaishin0527** 1 year, 11 months ago
Selected Answer: B
B: IMDSv2 is only fully supported on Nitro-based instances. If you need to enforce IMDSv2 on an instance, you must ensure that the instance is Nitro-based. These instances are designed to provide enhanced security, networking, and performance.
upvoted 3 times

**[Removed]** 1 year, 11 months ago
Selected Answer: C
Voting C
upvoted 2 times

**easytoo** 2 years ago
b-b-b-b-b-b
upvoted 1 times

**mawsman** 2 years ago
Selected Answer: A
The company's current infrastructure already supports both v1 and v2 - meaning instances are already nitro - as per anttan's comments the kernel needs to be 7.45 to update without outages
upvoted 1 times

**SONALID** 1 year, 8 months ago

nitro only supports kernel higher than 745. So it is not possible to have existing infrastructure in nitro without having kernel higher than 745

upvoted 1 times

**schalke04** 2 years, 4 months ago

**Selected Answer: B**

https://me.sap.com/notes/1656250

upvoted 3 times

**Kiran1982** 2 years, 4 months ago

**Selected Answer: B**

As described in note 1656250, IMDS version 2 is not supported with SAP in Xen based instances.

upvoted 4 times

**kk8s** 2 years, 4 months ago

C maybe

upvoted 2 times

Question #41 *Topic 1*

A company is running an SAP HANA database on AWS. The company wants to manage historical, infrequently accessed warm data for a native SAP HANA use case. An SAP solutions architect needs to recommend a solution that can provide online data storage in extended store, available for queries and updates. The solution must be an integrated component of the SAP HANA database and must allow the storage of up to five times more data in the warm tier than in the hot tier.
Which solution will meet these requirements?

A. Use Amazon Data Lifecycle Manager (Amazon DLM) with SAP Data Hub to move data in and out of the SAP HANA database to Amazon S3.

B. Use an SAP HANA extension node.

C. Use SAP HANA dynamic tiering as an optional add-on to the SAP HANA database.

D. Use Amazon Data Lifecycle Manager (Amazon DLM) with SAP HANA spark controller so that SAP HANA can access the data through the Spark SQL SDA adapter.

**Suggested Answer:** *A*

*Community vote distribution*

C (100%)

**schalke04** Highly Voted 2 years, 4 months ago

Selected Answer: C

c:
SAP HANA dynamic tiering is an integrated component of the SAP HANA database that allows the storage of warm data in an extended store. This solution enables the storage of up to five times more data in the warm tier compared to the hot tier. Dynamic tiering is available as an optional add-on for SAP HANA and provides online data storage in the extended store, making it available for queries and updates. This solution meets the requirement for an integrated component of the SAP HANA database and provides the ability to manage historical, infrequently accessed warm data for a native SAP HANA use case.

upvoted 6 times

**DigvijayGhosh** Most Recent 1 year, 7 months ago

Selected Answer: C

Option C

upvoted 1 times

**[Removed]** 1 year, 11 months ago

Selected Answer: C

option C

upvoted 1 times

**kk8s** 2 years, 5 months ago

Selected Answer: C

C - SAP HANA Dynamic Tiering
https://docs.aws.amazon.com/sap/latest/sap-hana/warm-data-tiering-options.html

upvoted 3 times

Question #42 *Topic 1*

A company plans to migrate its SAP NetWeaver deployment to AWS. The deployment runs on a Microsoft SQL Server database. The company plans to change the source database from SQL Server to SAP HANA as part of this process.
Which migration tools or methods should an SAP solutions architect use to meet these requirements? (Choose two.)

A. SAP HANA classical migration

B. SAP HANA system replication

C. SAP Software Update Manager (SUM) Database Migration Option (DMO) with System Move

D. SAP HANA backup and restore

E. SAP homogeneous system copy

**Suggested Answer:** *BD*

*Community vote distribution*

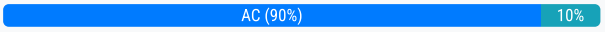


**DigvijayGhosh** 1 year, 7 months ago

**Selected Answer: AC**

A & C are correct

upvoted 1 times

**SONALID** 1 year, 8 months ago

Is Exam topics giving all incorrect answers purposefully? :D

upvoted 2 times

**SONALID** 1 year, 8 months ago

Correct ans - AC

upvoted 1 times

**kaishin0527** 1 year, 11 months ago

**Selected Answer: AC**

A,C:

A. SAP HANA classical migration: This is a two-step process where first you upgrade the existing system (if needed) and then perform the database migration. It's applicable when migrating from anyDB to SAP HANA.

C. SAP Software Update Manager (SUM) Database Migration Option (DMO) with System Move: DMO of the SUM is a tool that combines the upgrade and the migration of the SAP system to SAP HANA database into one process. In a one-step procedure, you can update an existing SAP system to a higher software version and migrate to SAP HANA.

upvoted 3 times

**[Removed]** 1 year, 11 months ago

**Selected Answer: BC**

Replication for HAHA
DMO for Netweaver

upvoted 1 times

**[Removed]** 1 year, 11 months ago

Edit, AC

upvoted 1 times

**easytoo** 2 years ago

a-c-a-c-a-c-a-c

upvoted 2 times

**sapien45** 2 years, 1 month ago

There is no such things as classicasl migration

upvoted 1 times

**[Removed]** 1 year, 11 months ago

There is https://docs.aws.amazon.com/sap/latest/sap-hana/migrating-hana-tools.html#migrating-hana-classical

upvoted 1 times

**Wedny** 2 years, 2 months ago

**Selected Answer: AC**

SAP provides tools and methodologies such as classical migration and SUM DMO to help its customers with the migration process for this scenario.

upvoted 1 times

**schalke04** 2 years, 4 months ago

**Selected Answer: AC**

A and C

upvoted 2 times

**kk8s** 2 years, 5 months ago

**Selected Answer: AC**

AC.

https://docs.aws.amazon.com/sap/latest/sap-hana/migrating-hana-anydb-to-hana.html

upvoted 2 times

Question #43 *Topic 1*

A company has an SAP Business One system that runs on SUSE Linux Enterprise Server 12 SP3. The company wants to migrate the system to AWS. An SAP solutions architect selects a homogeneous migration strategy that uses AWS Application Migration Service (CloudEndure Migration).
After the server migration process is finished, the SAP solutions architect launches an Amazon EC2 test instance from the R5 instance family. After a few minutes, the EC2 console reports that the test instance has failed an instance status check. Network connections to the instance are refused.
How can the SAP solutions architect solve this problem?

A. Reboot the instance to initiate instance migration to another host.

B. Request an instance limit increase for the AWS Region where the test instance is being launched.

C. Create a ticket for AWS Support that documents the test server instance ID. Wait for AWS to update the host of the R5 instance.

D. Install the missing drivers on the source system. Wait for the completion of migration synchronization. Launch the test instance again.

**Suggested Answer:** *D*

*Community vote distribution*


**kaishin0527** 1 year, 11 months ago

**Selected Answer: D**

D: This issue might occur if the source system lacks certain drivers required by the selected Amazon EC2 instance type. Before migrating, it's essential to ensure that the source system has all necessary drivers installed. If the drivers are missing, the server will not start correctly after the migration. Thus, the SAP solutions architect should install the missing drivers on the source system, wait for migration synchronization to complete, and then launch the test instance again.

upvoted 4 times

**[Removed]** 1 year, 11 months ago

**Selected Answer: A**

https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/monitoring-system-instance-status-check.html

Stopping and starting the instance can result in the instance being migrated to a new host

upvoted 1 times

**Peter290981** 2 years ago

D D D D D D D

upvoted 1 times

**easytoo** 2 years ago

a-a-a-a-a-a-a

upvoted 1 times

**mawsman** 2 years ago

**Selected Answer: D**

I think status checks would fail if the ENA driver wasn't installed so D

upvoted 2 times

**GiorgioGss** 2 years, 2 months ago

**Selected Answer: A**

https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/TroubleshootingInstances.html

upvoted 1 times

**SuiR** 2 years, 2 months ago

instance reboot will not trigger host change, only stop can change host, so A is incorrect.
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-lifecycle.html#lifecycle-differences

upvoted 1 times

**Sujit123kumar** 2 years, 2 months ago

**Selected Answer: D**

It should be D.

upvoted 1 times

**kumayuki** 2 years, 3 months ago

A

https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/TroubleshootingInstances.html

upvoted 1 times

**blanco750** 2 years, 3 months ago

This question doesn't make sense

upvoted 2 times

**matakuyy2** 2 years, 3 months ago

**Selected Answer: D**

I think B is different. If you are instance limit, it does not go to a status check.

Since this is a network connection problem, I think you are missing a network driver.

upvoted 1 times

**ohcn** 2 years, 4 months ago

I think is B

upvoted 2 times

**MKM** 2 years, 4 months ago

It's A

upvoted 1 times

**kk8s** 2 years, 5 months ago

i choose "B"

(maybe)

https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/troubleshooting-launch.html

upvoted 2 times

Question #44 *Topic 1*

An SAP basis architect is configuring high availability for a critical SAP system on AWS. The SAP basis architect is using an overlay IP address to route traffic to the subnets across multiple Availability Zones within an AWS Region for the system's SAP HANA database.
What should the SAP basis architect do to route the traffic to the Amazon EC2 instance of the active SAP HANA database?

A. Edit the route in the route table of the VPC that includes the EC2 instance that runs SAP HANSpecify the overlay IP address as the destination. Specify the private IP address of the EC2 instance as the target.

B. Edit the inbound and outbound rules in the security group of the EC2 instance that runs SAP HANA. Allow traffic for SAP HANA specific ports from the overlay IP address.

C. Edit the network ACL of the subnet that includes the EC2 instance that runs SAP HANA. Allow traffic for SAP HANA specific ports from the overlay IP address.

D. Edit the route in the route table of the VPC that includes the EC2 instance that runs SAP HANA. Specify the overlay IP address as the destination. Specify the elastic network interface of the EC2 instance as the target.

**Suggested Answer:** *B*

*Community vote distribution*


**SONALID** 1 year, 8 months ago

D
Add the overlay IP address in the Destination section and select Elastic Network Interface (ENI) name for one of your existing instances.
upvoted 1 times

**[Removed]** 1 year, 11 months ago

**Selected Answer: D**
Option D
upvoted 1 times

**juanvepe** 2 years ago

D.
Update routing tables
Add a routing entry to the routing tables which are assigned to the subnets of your primary and secondary EC2 instances. This IP address is the virtual IP (overlay IP) address of the SAP HANA cluster which needs to be outside the CIDR range of the VPC. To modify or add a route to a route table using the console:

Open the Amazon VPC console at https://console.aws.amazon.com/vpc/ (signin required).

In the navigation pane, choose Route Tables, and select the route table.

Choose Actions > Edit routes.

Scroll to the end of the list and click Add another route.

Add the overlay IP address in the Destination section and select Elastic Network Interface (ENI) name for one of your existing instances.

Save your changes by clicking Save routes.
upvoted 1 times

**easytoo** 2 years ago

d-d-d-d-d-d
upvoted 1 times

**MKM** 2 years, 4 months ago

Yes, the answer is D!
upvoted 4 times

**kk8s** 2 years, 5 months ago

**Selected Answer: D**

D i think.

https://docs.aws.amazon.com/sap/latest/sap-hana/sap-hana-on-aws-cluster-configuration-prerequisites.html

upvoted 3 times

Question #45 *Topic 1*

A company is running SAP ERP Central Component (SAP ECC) with a Microsoft SQL Server database on AWS. A solutions architect must attach an additional 1 TB Amazon Elastic Block Store (Amazon EBS) volume. The company needs to write the SQL Server database backups to this EBS volume before moving the database backups to Amazon S3 for long-term storage.
Which EBS volume type will meet these requirements MOST cost-effectively?

A. Throughput Optimized HDD (st1)

B. Provisioned IOPS SSD (io2)

C. General Purpose SSD (gp3)

D. Cold HDD (sc1)

**Suggested Answer:** *A*

*Community vote distribution*

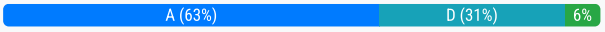


**trashi** Highly Voted 2 years, 4 months ago

Selected Answer: D

D
https://aws.amazon.com/ebs/cold-hdd/
upvoted 7 times

**Deepk12493** Highly Voted 2 years, 3 months ago

Selected Answer: A

https://docs.aws.amazon.com/sap/latest/sap-netweaver/backup-and-recovery.html
upvoted 5 times

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

Selected Answer: A

Answer is A.

D is not suitable because sc1 provides low throughput and high latency. The backup action could impact production
upvoted 1 times

**leotoras** 1 year, 5 months ago

SQL native tools to take backup on disk: Backup requires high throughput compared to IOPS. We recommend using Throughput Optimized HDD (st1) which provides maximum throughput of 500 MB/s per volume. Once the backup completes on disk, you can use scripts to move it to an Amazon S3 bucket
upvoted 2 times

**SONALID** 1 year, 8 months ago

Answer is A
SQL native tools to take backup on disk: Backup requires high throughput compared to IOPS. We recommend using Throughput Optimized HDD (st1) which provides maximum throughput of 500 MB/s per volume. Once the backup completes on disk, you can use scripts to move it to an Amazon S3 bucket.
upvoted 2 times

**G4Exams** 1 year, 9 months ago

A and D are possible but where does the requirement say to have high throughput ?! So more likely D that is more cost effective.
upvoted 2 times

**kaishin0527** 1 year, 11 months ago

Selected Answer: A

A: For workloads involving frequent, large, and sequential I/O operations, such as log processing and big data workloads, the st1 volume type is an economical and high-performance choice. Since database backups involve sequential write operations, Throughput Optimized HDD (st1) volumes, which are designed for large, sequential I/O workloads, will meet these requirements most cost-effectively.
upvoted 4 times

**[Removed]** 1 year, 11 months ago

Selected Answer: C

When unsure go for gp3

https://aws.amazon.com/blogs/storage/maximizing-microsoft-sql-server-performance-with-amazon-ebs/

upvoted 1 times

**[Removed]** 1 year, 11 months ago

Changing to A

upvoted 1 times

**easytoo** 2 years ago

a-a-a-a-a-a-a-a

upvoted 1 times

**mawsman** 2 years ago

Selected Answer: D

https://docs.aws.amazon.com/sap/latest/sap-hana/cold-data-tiering-options.html#sap-archiving

For archiving, another option is to use the Amazon Elastic Block Store (Amazon EBS) sc1 volume type as the underlying storage type for your archive file system. Amazon EBS sc1 volumes are inexpensive block storage and are designed for less frequently accessed workloads like data archiving. To increase durability and availability of your archived data, we recommend that you copy the data to Amazon S3 for backup and Amazon S3 Glacier for long term retention.

SC1 IS CHEAPEST

upvoted 3 times

**GiorgioGss** 2 years, 2 months ago

Selected Answer: A

"For SQL Server database backup, you can use one of the following methods:
SQL native tools to take backup on disk: Backup requires high throughput compared to IOPS. We recommend using Throughput Optimized HDD (st1) which provides maximum throughput of 500 MB/s per volume."

upvoted 1 times

**Kimzia** 2 years, 3 months ago

Selected Answer: A

(st1) volumes for SAP HANA to perform file-based backup. This volume type provides low-cost magnetic storage designed for large sequential workloads.

upvoted 2 times

**blanco750** 2 years, 3 months ago

Selected Answer: A

The question could be more clear mentioning throughput requirements but as other have mentioned, a couple of places mentions st1 as recommended HDD volume for backups as it gives high througput and for SAP backups high throughput is required.

upvoted 2 times

**schalke04** 2 years, 4 months ago

Selected Answer: A

A is correct.

upvoted 3 times

**trashi** 2 years, 4 months ago

Why A? sc1 is cheaper as said in https://aws.amazon.com/ebs/cold-hdd/ and a supported option for storing backups.

upvoted 3 times

**Hyperdanny** 2 years, 4 months ago

I am voting for D. This is not a Hana scenario though, but I don't see a reason to use throughput optimized data for backups in this scenario. Cold is cheaper.

upvoted 3 times

**Nots** 2 years, 4 months ago

Selected Answer: A

I'm torn between A and D.
There is also a description that st1 is an option when backing up to EBS as a local backup.
https://docs.aws.amazon.com/wellarchitected/latest/sap-lens/best-practice-14-2.html

upvoted 4 times

**forexamweb** 2 years, 4 months ago

**Selected Answer: D**

D maybe

https://docs.aws.amazon.com/sap/latest/sap-hana/cold-data-tiering-options.html#sap-archiving

upvoted 1 times

Question #46 *Topic 1*

Business users are reporting timeouts during periods of peak query activity on an enterprise SAP HANA data mart. An SAP system administrator has discovered that at peak volume, the CPU utilization increases rapidly to 100% for extended periods on the x1.32xlarge Amazon EC2 instance where the database is installed. However, the SAP HANA database is occupying only 1,120 GiB of the available 1,952 GiB on the instance. I/O wait times are not increasing. Extensive query tuning and system tuning have not resolved this performance problem.
Which solutions should the SAP system administrator use to improve the performance? (Choose two.)

A. Reduce the global_allocation_limit parameter to 1,120 GiB.

B. Migrate the SAP HANA database to an EC2 High Memory instance with a larger number of available vCPUs.

C. Move to a scale-out architecture for SAP HANA with at least three x1. 16xlarge instances.

D. Modify the Amazon Elastic Block Store (Amazon EBS) volume type from General Purpose to Provisioned IOPS for all SAP HANA data volumes.

E. Change to a supported compute optimized instance type for SAP HANA.

**Suggested Answer:** *DE*

*Community vote distribution*

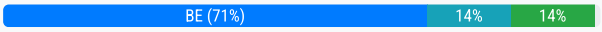


**koki2847** 1 year, 3 months ago

**Selected Answer: BC**

SAP HANA does not certify compute optimized instances. So E is not preferable I guess. https://docs.aws.amazon.com/sap/latest/general/sap-hana-aws-ec2.html

upvoted 1 times

**koki2847** 1 year, 3 months ago

"For SAP HANA databases that run in memory, memory optimized (r*, x*, u*) are your only options."
https://docs.aws.amazon.com/wellarchitected/latest/sap-lens/best-practice-13-2.html

upvoted 1 times

**acethetest1000** 1 year, 8 months ago

**Selected Answer: BE**

I think it relates to SAP HANA best practices as it recommends scaling up as much as possible before scaling out.

As AWS works based on t-shirt sizes and HANA demands memory optimized the next step is High Memory instance type, which in turn will increase the CPU amount.

upvoted 2 times

**G4Exams** 1 year, 9 months ago

**Selected Answer: BE**

B and E.

upvoted 1 times

**kaishin0527** 1 year, 11 months ago

**Selected Answer: BE**

B,E: This problem is about high CPU usage. The SAP HANA database is not fully utilizing the available memory, but CPU utilization is reaching 100% during peak query times. This suggests that the workload is CPU-bound. Therefore, you can alleviate this issue by adding more CPU resources, either by moving to a larger High Memory instance or switching to a compute-optimized instance type that has more vCPUs available.

upvoted 1 times

**[Removed]** 1 year, 11 months ago

**Selected Answer: BE**

Options B,E

upvoted 1 times

**[Removed]** 1 year, 11 months ago

Focus on the probable root cause of the performance issue, which is high CPU utilization.

upvoted 1 times

**easytoo** 2 years ago

b-c-b-c-b-c

upvoted 2 times

**Shaktimaan** 2 years, 3 months ago

A - WRONG, since no need to reduce memory.

B - WRONG, there is no ec2 instance larger than x1.32xlarge in X1 family.

C. RIGHT, Will add additional 1.5 times more CPU and Memory.

D. WRONG, No I/O issue as mentioned.

E. RIGHT, best way to get more compute and save on memory.

upvoted 2 times

**acethetest1000** 1 year, 8 months ago

The option B reads: move a High Memory instance which starts with 224vCPU and 3TB of memory: https://aws.amazon.com/ec2/instance-types/#memory-optimized

upvoted 1 times

**blanco750** 2 years, 3 months ago

Selected Answer: CE

Not 100% sure but eliminating the wrong ones

A. Reduce the global_allocation_limit parameter to 1,120 GiB. WRONG. Memory has nothing to do with High CPU here so this won't help

B. Migrate the SAP HANA database to an EC2 High Memory instance with a larger number of available vCPUs. Wrong. we need CPU optimized not memory

C. Move to a scale-out architecture for SAP HANA with at least three x1. 16xlarge instances. Scaling out actually helps in cases of high CPU utilization

D. Modify the Amazon Elastic Block Store (Amazon EBS) volume type from General Purpose to Provisioned IOPS for all SAP HANA data volumes. Issue is not related to high DISK I/O as it clearly says I/O wait is not increasing

E. Change to a supported compute optimized instance type for SAP HANA. this is probabaly the most correct option. use CPU optimized instances

upvoted 1 times

**Hyperdanny** 2 years, 4 months ago

I am voting for B/C: Change to a more powerful instance type (High Memory type) or scale out .

upvoted 3 times

**kk8s** 2 years, 4 months ago

C,D

https://docs.aws.amazon.com/wellarchitected/latest/sap-lens/best-practice-16-5.html

upvoted 4 times

**matakuyy2** 2 years, 3 months ago

I don't think D is the answer.

I don't think D is the answer, because it says "I/O wait times are not increasing", so I don't think I/O is the problem.

upvoted 1 times

**everydaysmile** 2 years, 4 months ago

can't choice B. Because the documentation says "but when that is not an option (such as scaling up a database instance), have a process in place to do so manually."

So i choose C and D

upvoted 2 times

Question #47 *Topic 1*

A company is moving to the AWS Cloud gradually. The company has multiple SAP landscapes on VMware. The company already has sandbox, development, and QA systems on AWS. The company's production system is still running on premises. The company has 2 months to cut over the entire landscape to the AWS Cloud.
The company has adopted a hybrid architecture for the next 2 months and needs to synchronize its shared file systems between the landscapes. These shared file systems include /trans directory mounts, /software directory mounts, and third-party integration mounts. In the on-premises landscape, the company has NFS mounts between the servers. On the AWS infrastructure side, the company is using Amazon Elastic File System (Amazon EFS) to share the common files.
An SAP solutions architect needs to design a solution to schedule transfer of these shared files bidirectionally four times each day. The data transfer must be encrypted.
Which solution will meet these requirements?

A. Write an rsync script. Schedule the script through cron for four times each day in the on-premises VMware servers to transfer the data from on premises to AWS.

B. Install an AWS DataSync agent on the on-premises VMware platform. Use the DataSync endpoint to synchronize between the on-premises NFS server and Amazon EFS on AWS.

C. Order an AWS Snowcone device. Use the Snowcone device to transfer data between the on-premises servers and AWS.

D. Set up a separate AWS Direct Connect connection for synchronization between the on-premises servers and AWS.

**Suggested Answer:** *B*

*Community vote distribution*

B (100%)

**[Removed]** 1 year, 11 months ago

Selected Answer: B

B. The solution is to install an AWS DataSync agent on the on-premises VMware platform. Use the DataSync endpoint to synchronize between the on-premises NFS server and Amazon EFS on AWS.

AWS DataSync is an online data movement and discovery service that simplifies data migration and helps you quickly, easily, and securely move your file or object data to, from, and between AWS storage services. DataSync can copy data to and from Network File System (NFS) file servers, Server Message Block (SMB) file servers, Amazon S3 buckets, Amazon EFS file systems, and AWS Snowcone.

Rsync is a utility for efficiently transferring and synchronizing files between a computer and an external hard drive and across networked computers by comparing the modification times and sizes of files. It is not designed for transferring data between on-premises servers and AWS.

upvoted 3 times

**[Removed]** 1 year, 11 months ago

https://aws.amazon.com/datasync/faqs/

upvoted 1 times

**easytoo** 2 years ago

b-b-b-b-b-b-b

upvoted 1 times

**schalke04** 2 years, 4 months ago

Selected Answer: B

B: AWS DataSync is an online data transfer service that simplifies, automates, and accelerates moving data between storage systems and services.

upvoted 4 times

Question #48 *Topic 1*

A company is planning to move to AWS. The company wants to set up sandbox and test environments on AWS to perform proofs of concept (POCs). Development and production environments will remain on premises until the POCs are completed.
At the company's on-premises location, SAProuter is installed on the same server as SAP Solution Manager. The company uses SAP Solution Manager to monitor the entire landscape. The company uses SAProuter to connect to SAP Support. The on-premises SAP Solution Manager instance must monitor the performance and server metrics of the newly created POC systems on AWS. The existing SAProuter must be able to report any issues to SAP.
What should an SAP solutions architect do to set up this hybrid infrastructure MOST cost-effectively?

A. Install a new SAP Solution Manager instance and a new SAProuter instance in the AWS environment. Connect the POC systems to these new instances. Use these new instances in parallel with the on-premises SAP Solution Manager instance and the on-premises SAProuter instance.

B. Install a new SAP Solution Manager instance and a new SAProuter instance in the AWS environment. Install the Amazon CloudWatch agent on all on-premises instances. Push the monitoring data to the new SAP Solution Manager instance. Connect all on-premises systems and POC systems on AWS to the new SAP Solution Manager instance and the new SAProuter instance. Remove the on-premises SAP Solution Manager instance and the on-premises SAProuter instance. Use the new instances on AWS.

C. Use AWS Site-to-Site VPN to connect the on-premises network to the AWS environment. Connect the POC systems on AWS to the on-premises SAP Solution Manager instance and the on-premises SAProuter instance.

D. Add the POC systems on AWS to the existing SAP Transport Management System that is configured in the on-premises SAP systems.

**Suggested Answer:** *B*

*Community vote distribution*

C (100%)

**acethetest1000** 1 year, 5 months ago

**Selected Answer: C**

It reads: the existing SAP Router must remain, hence I don't believe SolMan/SAP Router should be migrated. Besides, the company will start paying for another instance on AWS.

upvoted 1 times

**acethetest1000** 1 year, 5 months ago

**Selected Answer: C**

C is the most cost effective in terms of infrastructure and services.

upvoted 1 times

**[Removed]** 1 year, 11 months ago

**Selected Answer: C**

Going C

upvoted 1 times

**blanco750** 2 years, 3 months ago

**Selected Answer: C**

C

https://docs.aws.amazon.com/sap/latest/general/overview-sap-planning.html#figure-4

upvoted 2 times

**schalke04** 2 years, 4 months ago

**Selected Answer: C**

C makes sense

upvoted 2 times

**Hyperdanny** 2 years, 4 months ago

I am voting C:

https://docs.aws.amazon.com/sap/latest/general/overview-router-hybrid.html

upvoted 3 times

Question #49 *Topic 1*

An SAP solutions architect is using AWS Systems Manager Distributor to install the AWS Data Provider for SAP on production SAP application servers and SAP HANA database servers. The SAP application servers and the SAP HANA database servers are running on Red Hat Enterprise Linux.
The SAP solutions architect chooses instances manually in Systems Manager Distributor and schedules installation. The installation fails with an access and authorization error related to Amazon CloudWatch and Amazon EC2 instances. There is no error related to AWS connectivity.
What should the SAP solutions architect do to resolve the error?

A. Install the CloudWatch agent on the servers before installing the AWS Data Provider for SAP.

B. Download the AWS Data Provider for SAP installation package from AWS Marketplace. Use an operating system super user to install the agent manually or through a script.

C. Create an IAM role. Attach the appropriate policy to the role. Attach the role to the appropriate EC2 instances.

D. Wait until Systems Manager Agent is fully installed and ready to use on the EC2 instances. Use Systems Manager Patch Manager to perform the installation.

**Suggested Answer:** *A*

*Community vote distribution*

C (100%)

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

c - https://docs.aws.amazon.com/sap/latest/general/data-provider-troubleshooting.html

upvoted 5 times

**G4Exams** Most Recent 1 year, 9 months ago

Selected Answer: C

Most likely C

upvoted 1 times

**kaishin0527** 1 year, 11 months ago

Selected Answer: C

C: The AWS Data Provider for SAP requires access to Amazon CloudWatch and Amazon EC2 instances. It retrieves this access by assuming an IAM role that is attached to the EC2 instances. If this role is missing or doesn't have the appropriate permissions, the AWS Data Provider for SAP installation will fail with an access and authorization error. To resolve this error, you must create an IAM role and attach the appropriate policy to it, then attach this role to the EC2 instances.

upvoted 1 times

**schalke04** 2 years, 4 months ago

Selected Answer: C

C correct

upvoted 4 times

**Hyperdanny** 2 years, 4 months ago

I am voting for C. The Data Provider doesn't seem to have the right authorizations available.

upvoted 3 times

Question #50 *Topic 1*

A company is running its SAP applications on Oracle Database. Oracle Database is hosted on physical servers that are running SUSE Linux Enterprise Server. Because of compliance requirements, the company cannot install any additional software on its on-premises database servers. The company needs to migrate the SAP landscape to AWS and must continue to use Oracle Database.
Which migration solution should the company use to meet these requirements?

A. AWS Server Migration Service (AWS SMS)

B. AWS Application Migration Service (CloudEndure Migration)

C. SAP Software Update Manager (SUM) Database Migration Option (DMO) with System Move

D. Oracle Database replication with Oracle Data Guard

**Suggested Answer:** *D*

*Community vote distribution*

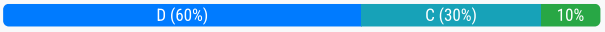


**MKM** Highly Voted 2 years, 4 months ago

The correct answer is D.
A) AWS SMS is for virtualized servers only
B) No additional software installation is allowed on-prem, so you cannot install CloudEndure Agent on-prem
C) SAP Software Update Manager (SUM) Database Migration Option (DMO) with System Move works only from AnyDB to SAP HANA migration (but here, the target system should be Oracle)
D) Yes.
This option would allow the company to continue using Oracle Database while meeting its compliance requirements, as it would not need to install additional software on its on-premises servers. They can use Oracle Data Guard to replicate their existing Oracle Database to an instance running on AWS, providing a way to migrate their SAP landscape to AWS while still using the same database.

upvoted 10 times

**schalke04** Highly Voted 2 years, 4 months ago

Selected Answer: D

D:
https://aws.amazon.com/blogs/database/physical-migration-of-oracle-databases-to-amazon-rds-custom-using-data-guard/

upvoted 5 times

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

I am confused because AWS does not support SUSE Linux for oracle

upvoted 1 times

**G4Exams** 1 year, 9 months ago

Selected Answer: D

D because the scenario is that no software so no agent etc can be installed ...

upvoted 1 times

**[Removed]** 1 year, 11 months ago

Selected Answer: C

Going for C

upvoted 1 times

**[Removed]** 1 year, 11 months ago

changing to D, MKM is right.

upvoted 1 times

**ADVIT** 1 year, 11 months ago

Selected Answer: C

It's C: https://docs.aws.amazon.com/sap/latest/sap-hana/migrating-hana-tools.html#migrating-hana-dmove2s4
Using SAP Database Migration Option (DMO) feature DMOVE2S4, you can migrate SAP ECC on SAP HANA or any other database, such as Oracle, SQL or others hosted on-premises to AWS Cloud.

upvoted 1 times

**GiorgioGss** 2 years, 2 months ago

**Selected Answer: A**

Based on the given requirements, the company needs to migrate the SAP landscape to AWS and continue to use Oracle Database without installing any additional software on the on-premises database servers.

Option C : (SAP Software Update Manager (SUM) Database Migration Option (DMO) with System Move) involves installing and using additional software (SUM) on the on-premises database servers, which violates the compliance requirements.

Option D : (Oracle Database replication with Oracle Data Guard) also requires additional software installation on the on-premises database servers.

Therefore, the best option is A (AWS Server Migration Service (AWS SMS)), which is a fully managed service that replicates on-premises servers, including their operating systems, applications, and data, to AWS without the need for installing any additional software on the on-premises servers. AWS SMS supports Oracle Database, and the company can continue to use it on AWS after migration.

upvoted 1 times

**trashi** 2 years, 4 months ago

**Selected Answer: C**

A) AWS SMS is for virtualized servers only

B) No additional software installation allowed on on-prem, so you cannot install CloudEndure Agent on-prem

C) Yes, will work fine

D) No additional software installation allowed on on-prem, so you cannot install Oracle Data Guard

upvoted 1 times

**SuiR** 2 years, 1 month ago

seems Oracle Data Guard is included in enterprise oracle version, so no need to install.

So D should be correct

upvoted 2 times

Question #51 *Topic 1*

A company is planning to migrate its SAP workloads to AWS. The company will use two VPCs. One VPC will be for production systems, and one VPC will be for non-production systems. The company will host the non-production systems and the primary node of all the production systems in the same Availability Zone.
What is the MOST cost-effective way to establish a connection between the production systems and the non-production systems?

A. Create an AWS Transit Gateway. Attach the VPCs to the transit gateway. Add the appropriate routes in the subnet route tables.

B. Establish a VPC peering connection between the two VPCs. Add the appropriate routes in the subnet route tables.

C. Create an internet gateway in each VPUse an AWS Site-to-Site VPN connection between the two VPCs. Add the appropriate routes in the subnet route tables.

D. Set up an AWS Direct Connect connection between the two VPCs. Add the appropriate routes in the subnet route tables.

**Suggested Answer:** *D*

*Community vote distribution*

B (100%)

**G4Exams** 1 year, 9 months ago

**Selected Answer: B**

If the VPCs are in the same AZ, and that is the case here, there will be no cost for traffic like if it goes via VPN through the internet. B is definitly the right answer here ..

upvoted 2 times

**kaishin0527** 1 year, 11 months ago

**Selected Answer: B**

B: VPC peering is a networking connection between two VPCs that enables you to route traffic between them using private IPv4 addresses or IPv6 addresses. It's a low-cost solution for interconnecting two VPCs, especially when they're located in the same AWS region.

upvoted 2 times

**[Removed]** 1 year, 11 months ago

**Selected Answer: B**

VPC peering if you only have 2 VPCs

upvoted 1 times

**easytoo** 2 years ago

Since the non-production systems and the primary node of the production systems are in the same Availability Zone, establishing a VPC peering connection is a straightforward and cost-effective solution. It does not incur any additional data transfer costs within the same AWS Region, making it an efficient choice.

Option A (AWS Transit Gateway) is a scalable and centralized solution, but it may introduce additional complexity and cost for this scenario, which involves only two VPCs. Option C (AWS Site-to-Site VPN) is typically used for secure connections between on-premises networks and AWS, and may not be necessary in this case where both VPCs are within AWS. Option D (AWS Direct Connect) is a dedicated network connection and may be overkill for connecting two VPCs within the same Availability Zone.

upvoted 2 times

**Azure1971** 2 years, 4 months ago

Answer is B:
Starting May 1st 2021, all data transfer over a VPC Peering connection that stays within an Availability Zone (AZ) is now free.

https://aws.amazon.com/about-aws/whats-new/2021/05/amazon-vpc-announces-pricing-change-for-vpc-peering/

upvoted 2 times

**schalke04** 2 years, 4 months ago

**Selected Answer: B**

B. Establish a VPC peering connection between the two VPCs. Add the appropriate routes in the subnet route tables.

upvoted 2 times

**MKM** 2 years, 4 months ago

The answer is B. Transit Gateway (Answer A) id not cost-effective compared to VPC peering.
upvoted 3 times

**kk8s** 2 years, 4 months ago
B for me
upvoted 2 times

**kk8s** 2 years, 4 months ago
As cost effective
upvoted 2 times

Question #52 *Topic 1*

An SAP engineer has deployed an SAP S/4HANA system on an Amazon EC2 instance that runs Linux. The SAP license key has been installed. After a while, the newly installed SAP instance presents an error that indicates that the SAP license key is not valid because the SAP system's hardware key changed. There have been no changes to the EC2 instance or its configuration.
Which solution will permanently resolve this issue?

A. Perform SAP kernel patching.

B. Apply a new SAP license that uses a new hardware key. Install the new key.

C. Set the SLIC_HW_VERSION Linux environment variable.

D. Reboot the EC2 instance.

**Suggested Answer:** *B*

*Community vote distribution*

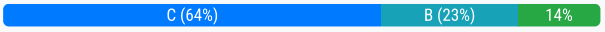


**tsangckl** 1 year, 2 months ago

Selected Answer: C

The SAP license key is tied to a hardware key, which is generated based on the system's hardware and OS configuration. In virtualized environments like Amazon EC2, some of these parameters might change, causing the hardware key to change and making the SAP license key invalid. By setting the SLIC_HW_VERSION environment variable, you can ensure that the hardware key remains constant even if the underlying hardware or OS configuration changes. This will prevent the SAP license key from becoming invalid in the future. Please note that you may need to restart the SAP system after setting this environment variable for the changes to take effect.

upvoted 1 times

**thuyeinaung** 1 year, 6 months ago

Selected Answer: C

I vote for C

upvoted 1 times

**SAPEXAMS** 1 year, 7 months ago

It is C as per the SAP Note 1178686
Extended algorithm (SLIC_HW_VERSION = 2)

upvoted 1 times

**G4Exams** 1 year, 9 months ago

Selected Answer: C

It is C

upvoted 3 times

**kaishin0527** 1 year, 11 months ago

Selected Answer: C

C: In certain cases, the hardware key can change even when there have been no changes to the EC2 instance or its configuration. This situation can cause the SAP license to become invalid. To prevent the hardware key from changing, you should set the Linux environment variable SLIC_HW_VERSION to the desired value. This value ensures a constant hardware key and prevents the SAP license from becoming invalid.

upvoted 3 times

**[Removed]** 1 year, 11 months ago

Selected Answer: B

A. WRONG because the newly install SAP instance updated the kernel (invalidating thus likely invalidating the license)
B. CORRECT, because the newly installed SAP instance generated a new hardware key, so a new key matching the new key must be installed
C. WRONG Setting the SLIC_HW_VERSION Linux environment variable will not resolve this issue
D. WRONG, Rebooting the EC2 instance will not resolve this issue.

B. Is correct in the context of the scenario in question assuming that the newly installed SAP instance updated the kernel.

RE: If a hardware key is generated before patching the SAP kernel installed license will invalid. This mean kernel patches will always invalidate a

license key.
https://docs.aws.amazon.com/sap/latest/general/overview-sap-on-aws.html
upvoted 1 times

**[Removed]** 1 year, 11 months ago

A. WRONG because the newly install SAP instance updated the kernel (invalidating thus likely invalidating the license)
B. CORRECT, because the newly installed SAP instance generated a new hardware key, so a new key matching the new key must be installed
C. WRONG Setting the SLIC_HW_VERSION Linux environment variable will not resolve this issue
D. WRONG, Rebooting the EC2 instance will not resolve this issue.

B. Is correct in the context of the scenario in question assuming that the newly installed SAP instance updated the kernel.

RE: If a hardware key is generated before patching the SAP kernel installed license will invalid. This mean kernel patches will always invalidate a license key.
https://docs.aws.amazon.com/sap/latest/general/overview-sap-on-aws.html
upvoted 1 times

**easytoo** 2 years ago

c-c-c-c-c-c
upvoted 1 times

**blanco750** 2 years, 3 months ago

Selected Answer: C

Check Appendix C in https://d0.awsstatic.com/enterprise-marketing/SAP/sap-on-aws-high-availability-guide.pdf
upvoted 2 times

**Azure1971** 2 years, 3 months ago

Answer is C:

SAP software is secured by a license using a hardware key (customer key). As described in SAP note 174911,"" this hardware key is based upon a MAC address of a network card in your Linux system"". In certain situations the algorithm can pick a different network card than it used for calculation before, thus returning a different hardware key.
SAP Note 1178686
upvoted 3 times

**anttan** 2 years, 4 months ago

Answer is B. The fact that the hardware key is regenerated is due to an update of the kernel. Since the kernel is already patched, there is no need to patch it again.
upvoted 1 times

**trashi** 2 years, 4 months ago

Selected Answer: C

A) is only a prerequisite, not a solution
B) is not a permanent solution as asked in the question
C) Yes, see SAP note 2184871
D) Rebooting the EC2 instance does not solve the license issue
upvoted 4 times

**schalke04** 2 years, 4 months ago

Selected Answer: B

B. Apply a new SAP license that uses a new hardware key. Install the new key.
upvoted 1 times

**trashi** 2 years, 4 months ago

Selected Answer: A

A because simply applying a new license will result in an invalid license later on again. As said in
https://docs.aws.amazon.com/sap/latest/general/overview-sap-on-aws.html
you need to apply a proper kernel patch before applying the license to PERMANENTLY solve the issue.
upvoted 3 times

**Balki** 2 years, 4 months ago

well said. patching is a prerequisite
upvoted 1 times

**kk8s** 2 years, 4 months ago

**Selected Answer: B**

B

https://docs.aws.amazon.com/sap/latest/general/overview-sap-on-aws.html

upvoted 3 times

Question #53 *Topic 1*

An SAP technology consultant needs to scale up a primary application server (PAS) instance. The PAS currently runs on a c5a.xlarge Amazon EC2 instance. The SAP technology consultant needs to change the instance type to c5a.2xlarge.
How can the SAP technology consultant meet this requirement?

A. Stop the complete SAP system. Stop the EC2 instance. Use the AWS Management Console or the AWS CLI to change the instance type. Start the EC2 instance. Start the complete SAP system.

B. While SAP is running, use the AWS Management Console or the AWS CLI to change the instance type without stopping the EC2 instance.

C. Stop the complete SAP system. Terminate the EC2 instance. Use the AWS Management Console or the AWS CLI to change the instance type. Start the EC2 instance. Start the complete SAP system.

D. While SAP is running, log in to the EC2 instance. Run the following AWS CLI command: aws ec2 modify-instance-attribute --instance-id <INSTANCEID> --instance-type "{\"Value\": \"c5a.2xlargel\"}".

**Suggested Answer:** *B*

*Community vote distribution*

A (100%)

**G4Exams** 1 year, 9 months ago

**Selected Answer: A**

Must be A. I dont know how you could ever change the instance tye while running.

upvoted 1 times

**kaishin0527** 1 year, 11 months ago

**Selected Answer: A**

A: To change the instance type in AWS, you need to stop the EC2 instance first. After stopping the instance, you can change its instance type using the AWS Management Console or the AWS CLI. Then, you can start the instance again. Keep in mind that all applications running on the instance should be stopped before stopping the instance itself.

upvoted 3 times

**[Removed]** 1 year, 11 months ago

**Selected Answer: A**

Voting option A

upvoted 1 times

**easytoo** 2 years ago

a-a-a-a-a-a-a-a

upvoted 1 times

**schalke04** 2 years, 4 months ago

**Selected Answer: A**

A. Stop the complete SAP system. Stop the EC2 instance. Use the AWS Management Console or the AWS CLI to change the instance type. Start the EC2 instance. Start the complete SAP system.

upvoted 2 times

**ohcn** 2 years, 4 months ago

A for me

upvoted 3 times

**Grillppl** 2 years, 4 months ago

I think A

upvoted 2 times

Question #54 *Topic 1*

A company has moved all of its SAP workloads to AWS. During peak business hours, end users are reporting performance issues because work processes are going into PRIV mode on an SAP S/4HANA system. An SAP support engineer indicates that SAP cannot provide support for this issue because some specific performance metrics are not available.
Which combination of actions must the company perform to comply with SAP support requirements? (Choose three.)

A. Buy an SAP license from AWS. Ensure that the SAP license is installed.

B. Select only an AWS Migration Acceleration Program (MAP) certified managed service provider (MSP).

C. Enable detailed monitoring for Amazon CloudWatch on each Amazon EC2 instance where SAP workloads are running.

D. Install, configure, and run the AWS Data Provider for SAP on each Amazon EC2 instance where SAP workloads are running.

E. Integrate AWS Systems Manager with SAP Solution Manager to provide alerts about SAP parameter configuration drift.

F. Enable SAP enhanced monitoring through a SAPOSCOL enhanced function.

**Suggested Answer:** *A*

*Community vote distribution*


**acethetest1000** 1 year, 5 months ago
CDF are the correct ones.
https://www.youtube.com/watch?v=mqWvwq1CNJI
upvoted 2 times

**Manojdeva** 1 year, 7 months ago
CDF is correct
upvoted 1 times

**Technolord** 1 year, 10 months ago
Selected Answer: C
CDF is the correct answer.
upvoted 2 times

**[Removed]** 1 year, 11 months ago
Selected Answer: C
voting CDF too
upvoted 1 times

**blanco750** 2 years, 3 months ago
Selected Answer: C
CDF is the answer
upvoted 3 times

**schalke04** 2 years, 4 months ago
Selected Answer: D
C D and F
upvoted 2 times

**Grillppl** 2 years, 4 months ago
I think CDF
upvoted 3 times

**ohcn** 2 years, 4 months ago
CDF - 1656250 - SAP on AWS: Support prerequisites
upvoted 4 times

Question #55 *Topic 1*

A company needs to implement high availability for its SAP S/4HANA system on AWS. The company will use a SUSE Linux Enterprise Server clustering solution in private subnets across two Availability Zones. An SAP solutions architect must ensure that the solution can route traffic to the active SAP instance in this clustered configuration.
What should the SAP solutions architect do to meet these requirements?

A. Implement the SAP cluster solution by using a secondary private IP address. Reassign the secondary private IP address from one network interface to another network interface in the event of any failure that affects the primary instance.

B. Implement the SAP cluster solution by using an Elastic IP address. Mask the failure of an instance or software by rapidly remapping the address to another instance in the account.

C. Implement the SAP cluster solution by using a public IP address. Use this public IP address for communication between the instances and the internet.

D. Implement the SAP cluster solution by using an overlay IP address that is outside the CIDR block of the VPC. Use overlay IP address routing to dynamically update the route table to point to the active node and provide external access by using a Network Load Balancer or AWS Transit Gateway.

**Suggested Answer:** *C*

*Community vote distribution*

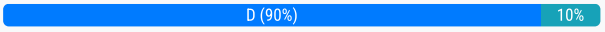


**schalke04** Highly Voted 2 years, 4 months ago

Selected Answer: D

The correct option is D. Implement the SAP cluster solution by using an overlay IP address that is outside the CIDR block of the VPC. Use overlay IP address routing to dynamically update the route table to point to the active node and provide external access by using a Network Load Balancer or AWS Transit Gateway.

upvoted 5 times

**acethetest1000** Most Recent 1 year, 5 months ago

Selected Answer: D

An overlay IP must be used.

upvoted 3 times

**DigvijayGhosh** 1 year, 7 months ago

Selected Answer: D

D is correct

upvoted 1 times

**kaishin0527** 1 year, 11 months ago

D: In a SUSE Linux Enterprise Server cluster for SAP, an overlay IP address is used for high availability. This overlay IP address is associated with the active node in the cluster and gets reassigned to another node in case of a failure. To make this work, it's necessary to dynamically update the routing table to point the overlay IP address to the correct, active instance.

upvoted 2 times

**[Removed]** 1 year, 11 months ago

Selected Answer: B

Voting B EIP + masking

upvoted 1 times

**[Removed]** 1 year, 11 months ago

Edit D.

upvoted 1 times

**ohcn** 2 years, 4 months ago

D - https://docs.aws.amazon.com/sap/latest/sap-netweaver/cluster-configuration-prereqs-sap-netweaver-ha.html#enable-cluster-instances-overlay-ip-sap-netweaver-ha

upvoted 4 times