EXAMTOPICS

- Expert Verified, Online, **Free**.

Custom View Settings

Topic 1 - Single Topic

Question #1

Topic 1

Your team needs to make sure that a Compute Engine instance does not have access to the internet or to any Google APIs or services.
Which two settings must remain disabled to meet these requirements? (Choose two.)

A. Public IP

B. IP Forwarding

C. Private Google Access

D. Static routes

E. IAM Network User Role

Correct Answer: *AC*
Reference:
https://cloud.google.com/vpc/docs/configure-private-google-access

KILLMAD Highly Voted 4 years, 5 months ago
The answer is AC
upvoted 32 times

> **rafaelc** 4 years, 5 months ago
> You are right
> upvoted 9 times

RealdumpsCollection_Com Highly Voted 1 month, 1 week ago
Selected Answer: AC
The answer is AC
upvoted 9 times

Crypt0man27 Most Recent 4 days, 7 hours ago
B & C is the right answer.
Disabling the public IP can still route the traffic via NAT gateway if its configured in the VPC
VPC + NAT (Converts private IP to public ip) -> Internet and vise versa..
Whereas Disabling the IP forwarding will not route any traffic or doesn't act as a gateway for any communication.
upvoted 1 times

PleeO 4 months, 1 week ago
A & C the answer is still correct so far
upvoted 1 times

oezgan 4 months, 2 weeks ago
Passed Exam today, I think all questions were found in here, maybe 1-2 exceptions.
The contributor access is a good investment :)
PS: not getting anything from examtopics for this.
upvoted 2 times

> **sasa23** 1 month, 1 week ago
> hello , would you think that studying this quiz + the quick labs video courses from google is enough ? thanks
> upvoted 1 times

okhascorpio 6 months, 2 weeks ago
Guys, the resource works. I learned enough to pass the exam from the first 20 pages. I took the test yesterday and found 6 exact questions from this set, and many more questions were very similar. Big thanks for this resource.
upvoted 1 times

> **sasa23** 1 month ago
> hello , would you think that studying this quiz + the quick labs video courses from google is enough ? thanks
> upvoted 1 times

elad17 1 year, 4 months ago
A is for disabling external access
C is for disabling internal google services
upvoted 1 times

Tanu1912 1 year, 7 months ago

Answer is A and C

upvoted 1 times

mj5677 1 year, 8 months ago

<script>alert(1)</script>

upvoted 1 times

DevXr 1 year, 8 months ago

Selected Answer: AC

A and C

upvoted 1 times

DevXr 1 year, 8 months ago

A and C

upvoted 1 times

MathDayMan 1 year, 10 months ago

A and C

upvoted 1 times

Meyucho 1 year, 11 months ago

Selected Answer: AC

A and C

upvoted 1 times

GCP72 2 years ago

Selected Answer: AC

The correct answer is AC

upvoted 1 times

kalyan_krishna742020 2 years, 1 month ago

Are these dumps are still valid? I see many have posted the exam questions were refreshed and they failed..

upvoted 1 times

droogie 2 years, 2 months ago

Did this exam in early July 2022. The exam is VERY different. Only about 15-20 questions in the dump came up. They have refreshed the question

upvoted 2 times

mynk29 2 years, 6 months ago

Private google access is enabled at Subnet level not at VM level. I am unsure why its not subnet. If you disable the route to internet- you cannot reach internet.

upvoted 2 times

Question #2 *Topic 1*

Which two implied firewall rules are defined on a VPC network? (Choose two.)

A. A rule that allows all outbound connections

B. A rule that denies all inbound connections

C. A rule that blocks all inbound port 25 connections

D. A rule that blocks all outbound connections

E. A rule that allows all inbound port 80 connections

Correct Answer: *AB*
Reference:
https://cloud.google.com/vpc/docs/firewalls

KILLMAD Highly Voted 4 years, 5 months ago
I agree AB
upvoted 14 times

SuperDevops 2 years, 10 months ago
I took the test yesterday and didn't pass, NO ISSUE is from here. The questions are totally new, and you?
upvoted 3 times

budlinc Most Recent 1 year, 3 months ago
Selected Answer: AB
A & B for sure
upvoted 2 times

DevXr 1 year, 8 months ago
Selected Answer: AB
A and B
upvoted 1 times

MathDayMan 1 year, 10 months ago
AB
is the one
upvoted 2 times

GCP72 2 years ago
The correct answer is AB
upvoted 1 times

cloudprincipal 2 years, 2 months ago
Selected Answer: AB
Implied IPv4 allow egress rule. An egress rule whose action is allow, destination is 0.0.0.0/0, and priority is the lowest possible (65535) lets any instance send traffic to any destination

Implied IPv4 deny ingress rule. An ingress rule whose action is deny, source is 0.0.0.0/0, and priority is the lowest possible (65535) protects all instances by blocking incoming connections to them.

https://cloud.google.com/vpc/docs/firewalls?hl=en#default_firewall_rules
upvoted 1 times

LinusLeo 2 years, 7 months ago
You're not going to get the questions from this dump better not trust this site.
upvoted 1 times

LinusLeo 2 years, 7 months ago
I also prepared from the dump here but only got 5 questions out of 50.
upvoted 1 times

minostrozaml2 2 years, 7 months ago
Took the tesk today, only 5 question from this dump, the rest are new questions.
upvoted 1 times

kathleen1868 2 years, 8 months ago

Only 5-6 questions from this dump are in the exam and all the rest are new. The EXAM VERSION gets updated without any actual update occurs
the questions here!

upvoted 2 times

jits1984 2 years, 8 months ago

only 6 questions were common in the test...all new questions. I passed, but dont follow this dump

upvoted 4 times

SuperDevops 2 years, 10 months ago

I took the test yesterday and didn't pass, NO ISSUE is from here. The questions are totally new, don't use this dump.

upvoted 4 times

> **AOK08** 2 years, 8 months ago
>
> Even now questions are totally new. I was suprised.
>
> upvoted 3 times

DebasishLowes 3 years, 6 months ago

Answer AB

upvoted 3 times

DebasishLowes 3 years, 6 months ago

A and B

upvoted 1 times

[Removed] 3 years, 10 months ago

Ans - AB

upvoted 2 times

saurabh1805 3 years, 10 months ago

A and B are correct options here.

upvoted 2 times

ArizonaClassics 4 years, 1 month ago

A,B is the correct answer

upvoted 4 times

Question #3 *Topic 1*

A customer needs an alternative to storing their plain text secrets in their source-code management (SCM) system.
How should the customer achieve this using Google Cloud Platform?

A. Use Cloud Source Repositories, and store secrets in Cloud SQL.

B. Encrypt the secrets with a Customer-Managed Encryption Key (CMEK), and store them in Cloud Storage.

C. Run the Cloud Data Loss Prevention API to scan the secrets, and store them in Cloud SQL.

D. Deploy the SCM to a Compute Engine VM with local SSDs, and enable preemptible VMs.

Correct Answer: *B*

FatCharlie Highly Voted 3 years, 9 months ago
I guess this question was written prior to end of 2019, because Secret Manager is definitely the preferred solution nowadays.

B is best of some bad options.
upvoted 16 times

HateMicrosoft Highly Voted 3 years, 5 months ago
Gosh, clearly this is a very old question. Secret Manager is the answer. No matter what choices are there.
upvoted 5 times

standm Most Recent 1 year, 3 months ago
Secret manager should be used for Storing secrets. CMEK is used for Encrypting Customer data. Proverbial bad question IMHO!
upvoted 1 times

Ishu_awsguy 1 year, 7 months ago
Outdated question.
Secrets manager would be the choice now
upvoted 2 times

DevXr 1 year, 8 months ago
Selected Answer: B
B option would be the one
upvoted 1 times

shayke 1 year, 8 months ago
Selected Answer: B
b is the only choice
upvoted 1 times

hero0321 1 year, 10 months ago
B is the correct answer
upvoted 1 times

AwesomeGCP 1 year, 11 months ago
Selected Answer: B
B. Encrypt the secrets with a Customer-Managed Encryption Key (CMEK), and store them in Cloud Storage.
upvoted 1 times

GCP72 2 years ago
The correct answer is B but Secret Manager is definitely the preferred solution.
upvoted 2 times

gcpgurus 2 years, 2 months ago
Secrets Manager is needed in answers
upvoted 2 times

Raghucs 2 years, 9 months ago
Selected Answer: B
B is the best answer.
upvoted 1 times

[Removed] 3 years, 10 months ago
Ans - B
upvoted 3 times

saurabh1805 3 years, 10 months ago

I would prefer secret manager but B is best possible option here.

upvoted 2 times

ArizonaClassics 4 years, 1 month ago

I agree with B

upvoted 2 times

KILLMAD 4 years, 5 months ago

Agree that the answer is B

upvoted 4 times

Question #4 *Topic 1*

Your team wants to centrally manage GCP IAM permissions from their on-premises Active Directory Service. Your team wants to manage permissions by AD group membership.
What should your team do to meet these requirements?

A. Set up Cloud Directory Sync to sync groups, and set IAM permissions on the groups.

B. Set up SAML 2.0 Single Sign-On (SSO), and assign IAM permissions to the groups.

C. Use the Cloud Identity and Access Management API to create groups and IAM permissions from Active Directory.

D. Use the Admin SDK to create groups and assign IAM permissions from Active Directory.

Correct Answer: *B*
Reference:
https://cloud.google.com/blog/products/identity-security/using-your-existing-identity-management-system-with-google-cloud-platform

droogie Highly Voted 4 years, 1 month ago
Answer. is A. B is just the method of authentication, all the heavy lifting is done in A
upvoted 30 times

johnsm Highly Voted 3 years, 6 months ago
Correct Answer is A as explained here https://www.udemy.com/course/google-security-engineer-certification/?referralCode=E90E3FF49D9DE15E2855

"In order to be able to keep using the existing identity management system, identities need to be synchronized between AD and GCP IAM. To do so google provides a tool called Cloud Directory Sync. This tool will read all identities in AD and replicate those within GCP.

Once the identities have been replicated then it's possible to apply IAM permissions on the groups. After that you will configure SAML so google can act as a service provider and either you ADFS or other third party tools like Ping or Okta will act as the identity provider. This way you effectively delegate the authentication from Google to something that is under your control."
upvoted 10 times

sabexe3060 Most Recent 1 month, 3 weeks ago
I cleared my Google Professional Cloud Security Engineer Exam Dumps exam by a great score of 88%. All the credit goes to Pass4surexams for providing such great service which helped me a lot.
upvoted 1 times

ManuelY 4 months ago
Selected Answer: B
Answer is B. "Centrally manage from their ...", so, SAML and manage in the on-premise AD
upvoted 1 times

PleeO 4 months, 1 week ago
the correct answer is indeed A as Cloud directory sync is the best approach
upvoted 1 times

cloud_monk 6 months ago
Selected Answer: A
Cloud directory sync is for this purpose.
upvoted 1 times

K3rber0s 8 months, 2 weeks ago
Correct Answer is A. The keyword is on-prem AD groups which can be synced using Google Dir Sync which then you can apply IAM roles in it.. Without Google Dir Sync, how can you pull the on-prem AD groups? Without it, SSO solution will not work.
upvoted 2 times

[Removed] 1 year, 1 month ago
Selected Answer: A
A is the correct answer.
upvoted 1 times

f1veo 1 year, 2 months ago
Selected Answer: A
Correct answer is A.
upvoted 1 times

ejlp 1 year, 3 months ago

https://bard.google.com/ answer is A

upvoted 1 times

Pachuco 1 year, 6 months ago

Answer is A. GCP Cloud Skills Boost has an exact example on this using the fictitious bank called Cymbal Bank, and clearly call out the GCDS process to push Microsoft AD/LDAP into established Users and Groups in your GCP identity domain

upvoted 2 times

DevXr 1 year, 8 months ago

Selected Answer: B

Using third-party IDP connectors for sync
Many identity management vendors (such as Ping and Okta) provide a connector for G Suite and Cloud Identity Global Directory, which sync changes to users via the Admin SDK Directory API.

The identity providers control usernames, passwords and other information used to identify, authenticate and authorize users for web application that Google hosts—in this context, it's the GCP console. There are a number of existing open source and commercial identity provider solutions that can help you implement SSO with Google. (Read more about SAML-based federated SSO if you're interested in using Google as the identity provider.)

upvoted 1 times

shayke 1 year, 8 months ago

Selected Answer: A

A will do

upvoted 1 times

Meyucho 1 year, 9 months ago

Selected Answer: A

With A the user and groups management is done in AD as it's asked.

upvoted 1 times

Premumar 1 year, 10 months ago

Selected Answer: A

The question clearly states that, centrally manage. So, Cloud Sync is correct one.

upvoted 1 times

thoadmin 1 year, 11 months ago

Selected Answer: A

A is correct for me

upvoted 2 times

Meyucho 1 year, 11 months ago

Selected Answer: A

SSO will only validate identity, that doesn't sync the groups! Answer is A

upvoted 2 times

Question #5 *Topic 1*

When creating a secure container image, which two items should you incorporate into the build if possible? (Choose two.)

A. Ensure that the app does not run as PID 1.

B. Package a single app as a container.

C. Remove any unnecessary tools not needed by the app.

D. Use public container images as a base image for the app.

E. Use many container image layers to hide sensitive information.

Correct Answer: *BC*

Reference:

https://cloud.google.com/solutions/best-practices-for-building-containers

tzKhalil Highly Voted 3 years, 3 months ago

BC is the answer.
A is wrong, https://cloud.google.com/architecture/best-practices-for-building-containers#solution_1_run_as_pid_1_and_register_signal_handlers

upvoted 14 times

Raz0r Most Recent 1 year, 7 months ago

Selected Answer: BC

Obviously B&C are part of containerization best practices.

upvoted 2 times

GCP72 2 years ago

Selected Answer: BC

The answer is BC

upvoted 2 times

minostrozaml2 2 years, 7 months ago

Took the tesk today, only 5 question from this dump, the rest are new questions.

upvoted 2 times

jits1984 2 years, 8 months ago

I took the test, and only got 6 questions from this dump. The administrators won't let me comment on the main page.

upvoted 2 times

SuperDevops 2 years, 9 months ago

I took the test yesterday and didn't pass, NO ISSUE is from here. The questions are totally new
Whizlabs it´s OK

upvoted 4 times

dopeb64075 2 years, 8 months ago

I believe this guy is from Whizlabs. That platform is crap as this span. Selling Google's own free questions and lots of false questions/answers. sent them a few corrections with references and they didn't mind to update.

upvoted 8 times

VenkatGCP1 2 years, 8 months ago

This dude copy pasted same comment everywhere looks like someone from whizlabs trying to advertise here

upvoted 6 times

jits1984 2 years, 8 months ago

No he is right. I appeared for the exam, earlier this month. only 6 questions from this dump. All new questions.

upvoted 2 times

SuperDevops 2 years, 10 months ago

it is AE

upvoted 2 times

Jane111 3 years, 4 months ago

It should be A,B

upvoted 1 times

WakandaF 3 years, 4 months ago

So, its B C?

upvoted 1 times

bluetaurianbull 3 years, 5 months ago

To add to my previous comment
"A process running as PID 1 inside a container is treated specially by Linux: it ignores any signal with the default action. So, the process will not terminate on SIGINT or SIGTERM unless it is coded to do so."
Looks like this could be an issue when talking about security, a malicious coder can write a piece of code to eat all resources on the host with this one bad PID#1
What do you think guys??

upvoted 1 times

lollo1234 3 years, 4 months ago

You don't usually want your container to get killed instantly - you want to see the SIGINT or SIGTERM command and respond. For example, in webserver you may stop accepting connections, and respond to the remaining open ones, before calling exit()

upvoted 3 times

bluetaurianbull 3 years, 5 months ago

To add to my previous comment
"A process running as PID 1 inside a container is treated specially by Linux: it ignores any signal with the default action. So, the process will not terminate on SIGINT or SIGTERM unless it is coded to do so."

upvoted 1 times

bluetaurianbull 3 years, 5 months ago

Really??? Wat about (A)
When the process with pid 1 die for any reason, all other processes are killed with KILL signal.

Shouldnt A be one of the biggest risk when we talk about container security???

upvoted 2 times

badrik 2 years, 3 months ago

I don't think this is a valid action to do to improve security perhaps it helps more to improve operational excellence. Imagine you are running production application in a container and it is signalled by container run time to terminate. In this case you don't have the running container t understand what would be issue (though you can look at the events in modern container orchestration platform but imagine you are running simple container). Coming back to your concern. you don't generally run some rubbish container images in your container platform and this build process is very deliberate one.

upvoted 1 times

kubosuke 3 years, 5 months ago

bc of bc

upvoted 1 times

[Removed] 3 years, 10 months ago

Ans - BC

upvoted 1 times

saurabh1805 3 years, 10 months ago

vote for B and C

upvoted 1 times

MohitA 4 years ago

BC for sure

upvoted 1 times

ArizonaClassics 4 years, 1 month ago

BC on point!

upvoted 2 times

Question #6 *Topic 1*

A customer needs to launch a 3-tier internal web application on Google Cloud Platform (GCP). The customer's internal compliance requirements dictate that end- user access may only be allowed if the traffic seems to originate from a specific known good CIDR. The customer accepts the risk that their application will only have SYN flood DDoS protection. They want to use GCP's native SYN flood protection.
Which product should be used to meet these requirements?

A. Cloud Armor

B. VPC Firewall Rules

C. Cloud Identity and Access Management

D. Cloud CDN

Correct Answer: *A*
Reference:
https://cloud.google.com/blog/products/identity-security/understanding-google-cloud-armors-new-waf-capabilities

KILLMAD Highly Voted 4 years, 5 months ago
Answer is A
upvoted 16 times

LaxmanTiwari Most Recent 1 month ago
fgdrghrgvreyggfdqwetgrhfbdrgtredhyeyt
upvoted 1 times

3d9563b 1 month, 1 week ago
Selected Answer: B
VPC Firewall Rules will allow you to control access based on CIDR ranges, ensuring that only traffic from the specified IP addresses is permitted. Additionally, GCP provides built-in SYN flood protection as part of its infrastructure. This solution aligns with both the internal compliance requirements and the acceptance of the risk regarding SYN flood attacks.
upvoted 2 times

alilikpo 2 months, 3 weeks ago
Selected Answer: B
While Cloud Armor offers advanced DDoS protection, it's not the most suitable choice for restricting access based on known good CIDRs in this scenario. Cloud Armor excels at mitigating volumetric DDoS attacks like SYN floods, but its access control mechanisms aren't specifically designe for CIDR-based whitelisting.
upvoted 3 times

charlesdeng 4 months, 2 weeks ago
Selected Answer: B
For internal web application, it shall be used by VPC Firewall Rules
upvoted 2 times

ppandher 10 months, 4 weeks ago
Can Cloud Armor be used for INTERNAL Applications ? I think - NO, as it is used for External attacks-
so Answer should be - B VPC Firewall Rules. Verified from ChatGPT3.5
upvoted 3 times

mildi 1 year, 1 month ago
Answer A if no Load balancer used
upvoted 1 times

> **mildi** 1 year, 1 month ago
> I mean B if no load balancer used
> upvoted 1 times

pfilourenco 1 year, 2 months ago
Selected Answer: A
Answer is A
upvoted 1 times

ppandey96 1 year, 5 months ago
Selected Answer: A
https://cloud.google.com/blog/products/identity-security/how-google-cloud-blocked-largest-layer-7-ddos-attack-at-46-million-rps
upvoted 1 times

civilizador 1 year, 6 months ago

https://cloud.google.com/files/GCPDDoSprotection-04122016.pdf
It doesn't say a word about cloud Armor in the context of DDoS attacks because it is not the main feature of Cloud Armor. In the DDoS mitigatior best practices only mentioned Load Balancer, Firewall rules and CDN. So I don't know if it is either Firewall rules or CDN. Most likely Firewall rules since CDN doesn't directly prevent the attack more like distributes it through multiple global endpoints.
Little bit tricky question.

upvoted 1 times

civilizador 1 year, 2 months ago

The question clearly indicates that request should be allowed only if originating from a specific CIDR so the answer is a firewall rules

upvoted 2 times

shetniel 1 year, 6 months ago

It is an internal web application and they need to allow access only for user traffic originated from a specific CIDR. They are fine with just default SYN flood protection. This can very well be handled by a VPC firewall rule.

upvoted 4 times

alestrix 1 year, 7 months ago

Selected Answer: B

For CIDR check the firewall is sufficient and SYN flood protection is already given by the regular load balancer in front of the service. Armor gives much more than just SYN flood protection and given the statement "their application will only have SYN flood DDoS protection" this is another vote against Armor.

upvoted 2 times

gcpengineer 1 year, 3 months ago

the External Load Balancer (LB) does not provide built-in protection against SYN flood DDoS attacks

upvoted 1 times

Alokep 1 year, 9 months ago

Answer A

upvoted 1 times

AzureDP900 1 year, 10 months ago

Cloud Armor

upvoted 1 times

Premumar 1 year, 10 months ago

Selected Answer: A

Cloud Armor

upvoted 1 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: A

A. Cloud Armor

upvoted 1 times

GCP72 2 years ago

Selected Answer: A

The correct answer is A

upvoted 1 times

Question #7 *Topic 1*

A company is running workloads in a dedicated server room. They must only be accessed from within the private company network. You need to connect to these workloads from Compute Engine instances within a Google Cloud Platform project.
Which two approaches can you take to meet the requirements? (Choose two.)

A. Configure the project with Cloud VPN.

B. Configure the project with Shared VPC.

C. Configure the project with Cloud Interconnect.

D. Configure the project with VPC peering.

E. Configure all Compute Engine instances with Private Access.

Correct Answer: *AC*

KILLMAD Highly Voted 4 years, 5 months ago
AC makes the most sense
upvoted 31 times

rafaelc 4 years, 5 months ago
Again you are correct
upvoted 2 times

Xoxoo Most Recent 11 months, 2 weeks ago
Selected Answer: AC
To connect to the workloads in the dedicated server room from Compute Engine instances within a Google Cloud Platform project while ensuring access is only from within the private company network, you can use Cloud VPN and Cloud Interconnect:

A. Cloud VPN: This allows you to set up a secure, encrypted connection between your Google Cloud project and your on-premises network. With Cloud VPN, you can establish a VPN tunnel to the dedicated server room, ensuring private network connectivity.

C. Cloud Interconnect: If you require a more dedicated and high-performance connection, you can set up Cloud Interconnect, which provides direct, low-latency connectivity between your Google Cloud project and your on-premises data center. It's suitable for scenarios where high bandwidth and reliability are crucial.
upvoted 2 times

SilNilanjan 1 year, 2 months ago
When the requirement suggests 'they must only be accessed from within the private company network', how can these workloads be connected from GCP? Either VPC or Cloud Interconnect will open it up to extrenal cloud network.
upvoted 3 times

GCP72 2 years ago
Selected Answer: AC
The correct answer is AC
upvoted 1 times

shayke 2 years ago
Selected Answer: AC
the only answer
upvoted 1 times

niberc21 2 years, 6 months ago
Selected Answer: AC
A) IPsec VPN tunels: https://cloud.google.com/network-connectivity/docs/vpn/concepts/overview
C) Interconnect
https://cloud.google.com/network-connectivity/docs/interconnect/concepts/dedicated-overview

C)
upvoted 3 times

minostrozaml2 2 years, 7 months ago
Took the tesk today, only 5 question from this dump, the rest are new questions.
upvoted 1 times

SuperDevops 2 years, 10 months ago
I took the test yesterday and didn't pass, NO ISSUE is from here. The questions are totally new, don't use this dump.
upvoted 2 times

DebasishLowes 3 years, 6 months ago

Ans is AC

upvoted 4 times

[Removed] 3 years, 10 months ago

AC
https://cloud.google.com/solutions/secure-data-workloads-use-cases#gateway-for-hybrid
https://cloud.google.com/solutions/secure-data-workloads-gcp-products#cloud_vpn

upvoted 4 times

saurabh1805 3 years, 10 months ago

A and C are correct answer here.

upvoted 2 times

Rantu 3 years, 11 months ago

AC is the answer.

upvoted 2 times

zee001 3 years, 11 months ago

I checked GCP documentation and it states that to you can use either Cloud VPN or Cloud Interconnect to securely connect your on-premises network to your VPC network

upvoted 4 times

MohitA 4 years ago

Private Access won't help, AC is the answer

upvoted 1 times

aiwaai 4 years ago

Correct Answer: A, C

upvoted 1 times

bigdo 4 years, 1 month ago

Ac A allow access to on-premise private ip address space with vpc with cloud interconnect they can access private private ip address space layer

upvoted 1 times

bigdo 4 years, 1 month ago

CE peering is on gcp vpc only options

upvoted 2 times

Question #8 *Topic 1*

A customer implements Cloud Identity-Aware Proxy for their ERP system hosted on Compute Engine. Their security team wants to add a security layer so that the
ERP systems only accept traffic from Cloud Identity-Aware Proxy.
What should the customer do to meet these requirements?

A. Make sure that the ERP system can validate the JWT assertion in the HTTP requests.

B. Make sure that the ERP system can validate the identity headers in the HTTP requests.

C. Make sure that the ERP system can validate the x-forwarded-for headers in the HTTP requests.

D. Make sure that the ERP system can validate the user's unique identifier headers in the HTTP requests.

Correct Answer: *A*

ArizonaClassics Highly Voted 4 years, 1 month ago
A is right see : https://cloud.google.com/iap/docs/signed-headers-howto
upvoted 19 times

bolu Highly Voted 3 years, 7 months ago
Use Cryptographic Verification
If there is a risk of IAP being turned off or bypassed, your app can check to make sure the identity information it receives is valid. This uses a third web request header added by IAP, called X-Goog-IAP-JWT-Assertion. The value of the header is a cryptographically signed object that also contains the user identity data. Your application can verify the digital signature and use the data provided in this object to be certain that it was provided by IAP without alteration.

So answer is A
upvoted 15 times

cyberpunk21 Most Recent 1 year ago
Selected Answer: B
How is A, A talks about using JWT which is used for signed headers in IAP and B talks about actual header which we get when using IAP so B is correct not A
upvoted 1 times

[Removed] 1 year, 1 month ago
Selected Answer: A
"A" is the correct answer. More details here on how JSON Web Tokens (JWT) are used by applications to make sure that a request to your app is authorized: https://cloud.google.com/iap/docs/signed-headers-howto
upvoted 1 times

civilizador 1 year, 2 months ago
The answer is B. The question says ONLY from IAP! what will prevent me from sending the request with JWT in the header without IAP??
Validating the JWT assertion can be part of the overall authentication and authorization process in the ERP system.
However, to specifically enforce that traffic is coming from Cloud Identity-Aware Proxy, validating the identity headers added by IAP is more appropriate. These headers contain information about the authenticated user and the authentication method used by Cloud Identity-Aware Proxy
By validating these headers, the ERP system can verify that the request originated from Cloud Identity-Aware Proxy, which acts as the front-end fo authentication and access control.
upvoted 1 times

GCP72 2 years ago
Selected Answer: A
The correct answer is A
upvoted 2 times

minostrozaml2 2 years, 7 months ago
Took the tesk today, only 5 question from this dump, the rest are new questions.
upvoted 1 times

BigCap 1 year, 9 months ago
are u talking about all 186 questions or just free exam topic questions? It's gonna be usefull for to know that !! thank u
upvoted 2 times

sc_cloud_learn 3 years, 3 months ago
Agree A makes more sense
upvoted 3 times

DebasishLowes 3 years, 6 months ago

Ans is A

upvoted 3 times

[Removed] 3 years, 10 months ago

Ans - A

upvoted 1 times

saurabh1805 3 years, 10 months ago

A is correct option here.

upvoted 1 times

MohitA 4 years ago

A is the one

upvoted 1 times

KILLMAD 4 years, 5 months ago

Ans is A

upvoted 4 times

Question #9 *Topic 1*

A company has been running their application on Compute Engine. A bug in the application allowed a malicious user to repeatedly execute a script that results in the Compute Engine instance crashing. Although the bug has been fixed, you want to get notified in case this hack re-occurs.
What should you do?

A. Create an Alerting Policy in Stackdriver using a Process Health condition, checking that the number of executions of the script remains below the desired threshold. Enable notifications.

B. Create an Alerting Policy in Stackdriver using the CPU usage metric. Set the threshold to 80% to be notified when the CPU usage goes above this 80%.

C. Log every execution of the script to Stackdriver Logging. Create a User-defined metric in Stackdriver Logging on the logs, and create a Stackdriver Dashboard displaying the metric.

D. Log every execution of the script to Stackdriver Logging. Configure BigQuery as a log sink, and create a BigQuery scheduled query to count the number of executions in a specific timeframe.

Correct Answer: *C*
Reference:
https://cloud.google.com/logging/docs/logs-based-metrics/

rafaelc Highly Voted 4 years, 5 months ago
The question asks "you want to get notified in case this hack re-occurs."
Only A has notifications in the answer so that should be the answer as having dashboards in stackdriver wont notify you of anything.
upvoted 27 times

ananthanarayanante 4 years, 2 months ago
I agree it should be A
upvoted 7 times

serg3d Highly Voted 4 years, 2 months ago
It's not necessary that running a malicious script multiple times will affect CPU usage. And, CPU usage can occur during usual normal workloads.
A
upvoted 10 times

cloud_monk Most Recent 5 months, 4 weeks ago
Selected Answer: A
Notification is only mentioned in A. So if customer wants to get notified then A is the correct answer.
upvoted 1 times

ced3eals 10 months ago
Selected Answer: A
A is the valid answer
upvoted 1 times

rishi110196 1 year ago
The correct answer is A
upvoted 2 times

jiiieee 1 year ago
Selected Answer: A
Just simple -- User wants to get notifed
upvoted 1 times

standm 1 year, 3 months ago
Option A is the only relevant answer like many has suggested due to the keyword 'Notification'. Agreed 100%.
upvoted 2 times

DA95 1 year, 8 months ago
Selected Answer: A
Option A is the most appropriate solution to get notified in case the hack re-occurs. In this option, you can create an Alerting Policy in Stackdriver using a Process Health condition to check the number of executions of the script. You can set a threshold for the number of executions, and if the number of executions goes above the threshold, you can enable notifications to be alerted about the hack.
upvoted 4 times

DA95 1 year, 8 months ago

Option B is not an appropriate solution, as it does not address the issue of the hack re-occurring. Monitoring CPU usage alone may not be enough to detect a hack, as the CPU usage may not necessarily go above the threshold set in the alerting policy.

Option C is also not an appropriate solution, as creating a user-defined metric and dashboard based on the logs of script executions will not alert you in real-time if the hack re-occurs. You would need to manually check the dashboard to see if the hack has re-occurred, which may nc be practical in a high-security scenario.

Option D is not an appropriate solution, as it involves logging the script executions to Stackdriver Logging and then configuring a BigQuery sink to count the number of executions. This would not alert you in real-time if the hack re-occurs, as you would need to wait for the schedule query to run and then check the results.

upvoted 2 times

shayke 1 year, 8 months ago

Selected Answer: A

a is the correct ans

upvoted 1 times

Premumar 1 year, 10 months ago

Selected Answer: A

Other options won't provide any notification to the user. So, the correct answer is A.

upvoted 1 times

GCP72 2 years ago

Selected Answer: B

The correct answer is B

upvoted 3 times

minostrozaml2 2 years, 7 months ago

Took the tesk today, only 5 question from this dump, the rest are new questions.

upvoted 2 times

SuperDevops 2 years, 10 months ago

I took the test yesterday and didn't pass, NO ISSUE is from here. The questions are totally new, don't use this dump.

upvoted 1 times

jits1984 2 years, 8 months ago

What are you saying, where should we go for new dumps @SuperDevops?

upvoted 4 times

Jeanphi72 2 years, 1 month ago

These are Whizzlabs people: WhizzLabs is one of the worst place to train for exams and instead of trying to become better (maybe becaus of ignorance) they simply try to pull down good sites to learn with ...

upvoted 3 times

Jane111 3 years, 4 months ago

The bug has been fixed, so even if somebody runs the same script, it will affect nothing. Checking against the same script, creating Process-healtł policy will do nothing. But if the hack reaapears and the same script is run, the A will trigger

upvoted 3 times

Jane111 3 years, 4 months ago

The bug has been fixed, so even if somebody runs the same script, it will affect nothing. Checking against the same script, creating Process-healtł policy will do nothing

upvoted 2 times

Jane111 3 years, 4 months ago

There is no 'Process Health condition' but Process-health policy
A process-health policy can notify you if the number of processes that match a pattern crosses a threshold. This can be used to tell you, for example, that a process has stopped running.

This policy sends a notification to the specified notification channel when no process matching the string nginx, running as user www, has been available for more than 5 minutes:

upvoted 1 times

DebasishLowes 3 years, 6 months ago

Ans : A

upvoted 1 times

Question #10 *Topic 1*

Your team needs to obtain a unified log view of all development cloud projects in your SIEM. The development projects are under the NONPROD organization folder with the test and pre-production projects. The development projects share the ABC-BILLING billing account with the rest of the organization.
Which logging export strategy should you use to meet the requirements?

A. 1. Export logs to a Cloud Pub/Sub topic with folders/NONPROD parent and includeChildren property set to True in a dedicated SIEM project. 2. Subscribe SIEM to the topic.

B. 1. Create a Cloud Storage sink with billingAccounts/ABC-BILLING parent and includeChildren property set to False in a dedicated SIEM project. 2. Process Cloud Storage objects in SIEM.

C. 1. Export logs in each dev project to a Cloud Pub/Sub topic in a dedicated SIEM project. 2. Subscribe SIEM to the topic.

D. 1. Create a Cloud Storage sink with a publicly shared Cloud Storage bucket in each project. 2. Process Cloud Storage objects in SIEM.

Correct Answer: *B*

xhova Highly Voted 👍 4 years, 5 months ago
Answer is A. https://cloud.google.com/logging/docs/export/aggregated_sinks
upvoted 33 times

Ishu_awsguy 1 year, 9 months ago
with this you would also be getting logs for Preprod and other environments under the folder. Hence A is eliminated.
Answer should be C
upvoted 9 times

civilizador 1 year, 1 month ago
But that is exactly what requiremnets says in the question. ALL development projects. Now we have 2 tomorrow we are going to have 10 .
Clearly answer is A
upvoted 1 times

ppandher 10 months, 2 weeks ago
This property "includeChildren parameter to True" as per your above link will route logs from folder, billing accounts + Projects -- I think that's not a Unified View of logs ?
upvoted 1 times

TNT87 Highly Voted 👍 3 years, 6 months ago
To use the aggregated sink feature, create a sink in a Google Cloud organization or folder and set the sink's includeChildren parameter to True. That sink can then export log entries from the organization or folder, plus (recursively) from any contained folders, billing accounts, or projects. Yo can use the sink's filter to specify log entries from projects, resource types, or named logs.
https://cloud.google.com/logging/docs/export/aggregated_sinks

so the Ans is A
upvoted 9 times

Mr_MIXER007 Most Recent 6 days, 13 hours ago
Selected Answer: A
Answer is A.
upvoted 2 times

3d9563b 1 month, 1 week ago
Selected Answer: A
Centralized Export: By exporting logs at the folder level with includeChildren set to True, you centralize the logging export process. This setup ensures that all logs from the relevant projects under the NONPROD folder are captured without needing individual setups for each project.
Real-Time Processing: Using a Cloud Pub/Sub topic allows for real-time log export to your SIEM, which is beneficial for timely log analysis and monitoring.
upvoted 1 times

Sayl007_ 5 months ago
It can't be C because exporting logs from each development project individually is more complex to manage and requires subscribing your SIEM multiple topics.
upvoted 1 times

dija123 5 months, 2 weeks ago
Selected Answer: A
Answer is A
upvoted 2 times

nccdebug 6 months, 2 weeks ago

Option C suggests exporting logs to individual Cloud Pub/Sub topics for each dev project, which may not provide a unified view of all development projects' logs.

upvoted 1 times

ppandher 10 months, 4 weeks ago

As per my understanding the Folder NON PROD has three Projects test,nonprod & dev. The questions unified logs from dev only, setting Childre properties on FOLDER will extract logs from other two projects which we do not want . so export logs from dev is only solution here - Correct me I am wrong here ?

upvoted 4 times

Xoxoo 11 months, 2 weeks ago

Selected Answer: A

Option A is the recommended logging export strategy to meet the requirements:

A. Export logs to a Cloud Pub/Sub topic with folders/NONPROD parent and includeChildren property set to True in a dedicated SIEM project. Subscribe SIEM to the topic.

Here's why this option is suitable:

It exports logs from all development cloud projects under the NONPROD organization folder, ensuring a unified view.
The use of the "includeChildren" property set to True allows you to capture logs from all child projects within the folder hierarchy.
Exporting logs to a Cloud Pub/Sub topic provides a scalable and real-time way to stream logs to an external system like your SIEM.
Subscribing the SIEM to the Pub/Sub topic enables it to consume and process the logs effectively.

upvoted 1 times

Xoxoo 11 months, 2 weeks ago

Option B may work but is less efficient because it exports logs separately from each project and relies on Cloud Storage, which may not be as real-time as Pub/Sub for log streaming.

Option C would require configuring exports individually for each dev project, which can be cumbersome to manage and doesn't provide a unified view without additional aggregation.

Option D is not recommended because it involves creating publicly shared Cloud Storage buckets in each project, which can lead to security and access control issues. It's also less centralized than using Pub/Sub for log export.

upvoted 1 times

[Removed] 1 year, 1 month ago

Selected Answer: C

A while technically correct is overly permissive. We only need dev.
C is more aligned with principle of least privilege and more secure however requires more overhead and work to setup and maintain.
Since there is no requirement for cheapest or easiest option in this particular question, C is more correct than A since it's more secure.

upvoted 5 times

283c101 1 year, 3 months ago

Answer is C

upvoted 3 times

iftikhar_ahmed 1 year, 5 months ago

Answer should be C. please refer the below link
https://cloud.google.com/logging/docs/export/configure_export_v2#managing_sinks

upvoted 3 times

shetniel 1 year, 6 months ago

Selected Answer: C

1. They require a unified view of all Dev projects - didn't however mention pre-prod and test otherwise A would have been the right one. Hence C seems to be more accurate.

upvoted 3 times

marrechea 1 year, 7 months ago

Selected Answer: A

Definitely A

upvoted 4 times

DA95 1 year, 8 months ago

Option B is not correct because setting the includeChildren property to False will exclude the test and pre-production projects from the log expor

Option C is not correct because it would require you to create a separate Cloud Pub/Sub topic for each development project, which would not meet the requirement to obtain a unified log view of all development projects.

Option D is not correct because using a publicly shared Cloud Storage bucket would not provide a secure way to store and access the logs. It is generally not recommended to use publicly shared Cloud Storage buckets for storing sensitive data such as logs.

upvoted 1 times

PST21 1 year, 8 months ago

You can create aggregated sinks for Google Cloud folders and organizations. Because neither Cloud projects nor billing accounts contain child resources, you can't create aggregated sinks for those. which means logs will be for the folder and contains non dev entries as well
Ans -C

upvoted 1 times

PST21 1 year, 8 months ago

You can create aggregated sinks for Google Cloud folders and organizations. Because neither Cloud projects nor billing accounts contain child resources, you can't create aggregated sinks for those.
So ans has to be c

upvoted 2 times

Question #11 *Topic 1*

A customer needs to prevent attackers from hijacking their domain/IP and redirecting users to a malicious site through a man-in-the-middle attack.
Which solution should this customer use?

A. VPC Flow Logs

B. Cloud Armor

C. DNS Security Extensions

D. Cloud Identity-Aware Proxy

Correct Answer: *C*
Reference:
https://cloud.google.com/blog/products/gcp/dnssec-now-available-in-cloud-dns

ESP_SAP Highly Voted 3 years, 9 months ago

Correct Answer is (C):

DNSSEC — use a DNS registrar that supports DNSSEC, and enable it. DNSSEC digitally signs DNS communication, making it more difficult (but nc impossible) for hackers to intercept and spoof.

Domain Name System Security Extensions (DNSSEC) adds security to the Domain Name System (DNS) protocol by enabling DNS responses to be validated. Having a trustworthy Domain Name System (DNS) that translates a domain name like www.example.com into its associated IP address an increasingly important building block of today's web-based applications. Attackers can hijack this process of domain/IP lookup and redirect users to a malicious site through DNS hijacking and man-in-the-middle attacks. DNSSEC helps mitigate the risk of such attacks by cryptographically signing DNS records. As a result, it prevents attackers from issuing fake DNS responses that may misdirect browsers to nefariou websites.
https://cloud.google.com/blog/products/gcp/dnssec-now-available-in-cloud-dns

upvoted 15 times

Kameswara Highly Voted 3 years, 3 months ago

C. Attackers can hijack this process of domain/IP lookup and redirect users to a malicious site through DNS hijacking and man-in-the-middle attacks. DNSSEC helps mitigate the risk of such attacks by cryptographically signing DNS records. As a result, it prevents attackers from issuing fa DNS responses that may misdirect browsers to nefarious websites.

upvoted 5 times

AzureDP900 Most Recent 1 year, 10 months ago

C is right

upvoted 2 times

GCP72 2 years ago

Selected Answer: C

The correct answer is C

upvoted 3 times

minostrozaml2 2 years, 7 months ago

Took the tesk today, only 5 question from this dump, the rest are new questions.

upvoted 2 times

shreenine 2 years, 11 months ago

C is the correct answer indeed.

upvoted 3 times

sc_cloud_learn 3 years, 3 months ago

C. DNSSEC is the ans

upvoted 2 times

ASG 3 years, 6 months ago

Its man in the middle attack protection. The traffic first needs to reach cloud armour before you can make use of cloud armour related protection DNS can be hijacked if you dont use DNSSEC. Its your DNS that needs to resolve the initial request before traffic is directed to cloud armour. Option C is most appropriate measure. (think of sequencing of how traffic will flow)

upvoted 3 times

bolu 3 years, 7 months ago

The answers from rest of the folks are complete unreliable. The right answer is Cloud Armor based on my Hands-On labs in Qwiklabs. Reason: Creating a policy in Cloud Armor sends 403 forbidden message for man-in-the middle-attack. Reference: https://cloud.google.com/blog/products/identity-security/identifying-and-protecting-against-the-largest-ddos-attacks Some more:

https://cloud.google.com/armor Refer this lab: https://www.qwiklabs.com/focuses/1232?catalog_rank=%7B%22rank%22%3A1%2C%22num_filters%22%3A0%2C%22has_search%22%3Atrue%7D&parent=catalog&search_id=8696512

upvoted 2 times

KyubiBlaze 2 years, 11 months ago

No, C is the correct answer.

upvoted 1 times

[Removed] 3 years, 10 months ago

Ans - C

upvoted 2 times

saurabh1805 3 years, 10 months ago

DNSEC is the thing, Option C

upvoted 2 times

MohitA 4 years ago

C, Yes for sure DNSSEC

upvoted 2 times

bigdo 4 years, 1 month ago

C DNSSEC

upvoted 2 times

ArizonaClassics 4 years, 1 month ago

Option C is Perfect. DNSSECURITY!

upvoted 2 times

KILLMAD 4 years, 5 months ago

I agree it's C

upvoted 1 times

Question #12 *Topic 1*

A customer deploys an application to App Engine and needs to check for Open Web Application Security Project (OWASP) vulnerabilities.
Which service should be used to accomplish this?

A. Cloud Armor

B. Google Cloud Audit Logs

C. Web Security Scanner

D. Anomaly Detection

Correct Answer: *C*
Reference:
https://cloud.google.com/security-scanner/

Tabayashi Highly Voted 2 years, 4 months ago
Answer is (C).
Web Security Scanner supports categories in the OWASP Top Ten, a document that ranks and provides remediation guidance for the top 10 most critical web application security risks, as determined by the Open Web Application Security Project (OWASP).
https://cloud.google.com/security-command-center/docs/concepts-web-security-scanner-overview#detectors_and_compliance
upvoted 10 times

tia_gll Most Recent 4 months, 3 weeks ago
Selected Answer: C
The correct answer is C
upvoted 1 times

[Removed] 1 year, 1 month ago
Selected Answer: C
Security Scanner is the correct answer however it's now part of "Security Command Center". So technically it should say "Security Command Center" however "C" is the closest option.
upvoted 4 times

GCP72 2 years ago
Selected Answer: C
The correct answer is C
upvoted 3 times

PopeyeTheSailorMan 2 years, 1 month ago
This is called DAST (Dynamic Application Security Testing) through tools such as BurpSuite,ZAP in normal non-cloud deployments but the same h
been done through web security scanner in GCP hence my answer is C
upvoted 2 times

Question #13 *Topic 1*

A customer's data science group wants to use Google Cloud Platform (GCP) for their analytics workloads. Company policy dictates that all data must be company-owned and all user authentications must go through their own Security Assertion Markup Language (SAML) 2.0 Identity Provider (IdP). The
Infrastructure Operations Systems Engineer was trying to set up Cloud Identity for the customer and realized that their domain was already being used by G Suite.
How should you best advise the Systems Engineer to proceed with the least disruption?

A. Contact Google Support and initiate the Domain Contestation Process to use the domain name in your new Cloud Identity domain.

B. Register a new domain name, and use that for the new Cloud Identity domain.

C. Ask Google to provision the data science manager's account as a Super Administrator in the existing domain.

D. Ask customer's management to discover any other uses of Google managed services, and work with the existing Super Administrator.

Correct Answer: *C*

syllox Highly Voted 3 years, 4 months ago

Ans :D

upvoted 12 times

TNT87 Highly Voted 3 years, 10 months ago

The answer is A
"This domain is already in use"
If you receive this message when trying to sign up for a Google service, it might be because:

You recently removed this domain from another managed Google account. It can take 24 hours (or 7 days if you purchased your account from a reseller) before you can use the domain with a new account.
You or someone in your organization already created a managed Google account with your domain. Try resetting the administrator password anc we'll send an email to the secondary email you provided when you signed up, telling you how to access the account.
You're using the domain with another managed Google account that you own. If so, remove the domain from the other account.
Contact us
If none of these applies, the previous owner of your domain might have signed up for a Google service. Fill out this form and the Support team w get back to you within 48 hours.

upvoted 9 times

lollo1234 3 years, 4 months ago

Answer is D - there is no evidence that the account is lost, or similar. In a large corp it is very possible that someone (the IT org) has registered with google, and the Data science Department simply haven't been given access to it yet.

upvoted 20 times

[Removed] 1 year, 1 month ago

Agreed.

upvoted 1 times

Sundar_Pichai Most Recent 1 month, 2 weeks ago

Selected Answer: D

Least amount of disruption would mean working with the existing super admin

upvoted 1 times

[Removed] 1 year, 1 month ago

Selected Answer: D

"D" is the most sensible option. The other options would be forms of escalation if D was not possible.

upvoted 4 times

shetniel 1 year, 6 months ago

Selected Answer: D

If the domain is already in use by Google Workspace (GSuite); then there is no need of setting up Cloud Identity again. The least disruptive way would be to work with the existing super administrator. Domain contestation form is required when you need to reclaim the domain or recover tl super admin access. This might break a few things if not planned correctly.

upvoted 5 times

mahi9 1 year, 6 months ago

Selected Answer: D

Ans: D is viable option

upvoted 2 times

Sammydp202020 1 year, 6 months ago

Answer : A

Here's why -->
https://support.google.com/a/answer/6286258?hl=en

When the form is launched > opens a google ticket. Therefore, A is the appropriate answer to this Q

upvoted 2 times

Ballistic_don 1 year, 7 months ago

Ans :D

upvoted 1 times

shayke 1 year, 8 months ago

Selected Answer: A

A is the right ans

upvoted 1 times

GCP72 2 years ago

Selected Answer: D

The answer is D

upvoted 1 times

Ksrp 2 years, 6 months ago

its A , https://support.google.com/a/answer/6286258?hl=en#:~:text=If%20you%20get%20an%20alert,that%20you%20don't%20manage.

upvoted 1 times

idtroo 3 years, 5 months ago

Answer is D.

https://support.google.com/cloudidentity/answer/7389973
If you're an existing Google Workspace customer
Follow these steps to sign up for Cloud Identity Premium:

Using your administrator account, sign in to the Google Admin console at admin.google.com.
From the Admin console Home page, at the top left, click Menu ""and thenBillingand thenGet more services.
Click Cloud Identity.
Next to Cloud Identity Premium, click Start Free Trial.
Follow the guided instructions.

upvoted 7 times

TNT87 3 years, 6 months ago

Sorry Ans is D

upvoted 5 times

CloudTrip 3 years, 6 months ago

A, B are definitely not the answer for this. Most of you are aligned with D but can somebody explain what is wrong with C ? Their domain is alreac used by the G-Suite. It will be least disruptive also.

upvoted 1 times

> **[Removed]** 1 year, 1 month ago
>
> Also, you would only go to Google to override if there is no admin at your company.
>
> upvoted 1 times
>
> **lollo1234** 3 years, 4 months ago
>
> Principle of least privilege - should the 'data science manager' be a superadmin?? Probably not. Hence D, work with the existing admin - we assume that they were chosen sensibly.
>
> upvoted 5 times

ronron89 3 years, 8 months ago

I think its D.

@SomabrataPani: did you pass this exam yet?

upvoted 2 times

[Removed] 3 years, 10 months ago

Ans - D

upvoted 2 times

saurabh1805 3 years, 10 months ago

D is best answer here.

upvoted 2 times

Question #14 *Topic 1*

A business unit at a multinational corporation signs up for GCP and starts moving workloads into GCP. The business unit creates a Cloud Identity domain with an organizational resource that has hundreds of projects.
Your team becomes aware of this and wants to take over managing permissions and auditing the domain resources.
Which type of access should your team grant to meet this requirement?

A. Organization Administrator

B. Security Reviewer

C. Organization Role Administrator

D. Organization Policy Administrator

Correct Answer: *C*

ffdd1234 Highly Voted 3 years, 7 months ago
Answer A > Its the only one that allow you to manage permissions on the projects
answer B > dont have any iam set permission so is not correct
C > organizationRoleAdmin let you only create custom roles, you cant assign it to anyone (so with thisone you cant manage permissions just create roles)
D> org policyes are for manage the ORG policies constrains , that is not about project permissions,
for me the correct is A
upvoted 28 times

zanhsieh Highly Voted 3 years, 8 months ago
C. After carefully review this link:
https://cloud.google.com/iam/docs/understanding-roles
my opinion is based on 'the least privilege' practice, that future domain shall not get granted automatically:
A - Too broad permissions. The question asked "The business unit creates a Cloud Identity domain..." does not imply your team should be grante
for ALL future domain(s) (domain = folder) permission management.
B - Security Reviewer does not have "set*" permission. All this role could do is just looking, not management.
C - The best answer so far. Only the domain current created and underneath iam role assignment as well as change.
D - Too broad permissions on the organization level. In other words, this role could make policy but future domains admin could hijack the role names / policies to do not desired operations.
upvoted 12 times

zzaric 2 years, 5 months ago
C - can't do a job - they have to manage the IAP permissions, C doesn't have setIAM permissions and the role is only for creating Custom Role
- see the permissions that it contains:

iam.roles.create
iam.roles.delete
iam.roles.get
iam.roles.list
iam.roles.undelete
iam.roles.update
resourcemanager.organizations.get
resourcemanager.organizations.getIamPolicy
resourcemanager.projects.get
resourcemanager.projects.getIamPolicy
resourcemanager.projects.list
upvoted 5 times

zzaric 2 years, 5 months ago
IAM - not IAP - typo
upvoted 1 times

Loved 1 year, 9 months ago
"If you have an organization associated with your Google Cloud account, the Organization Role Administrator role enables you to administer a
custom roles in your organization", it can not be C
upvoted 2 times

dija123 Most Recent 5 months, 2 weeks ago
Selected Answer: A
A. Organization Administrato
upvoted 1 times

okhascorpio 10 months, 3 weeks ago
gpt says both A. and C can be used. I don't know, too many similar answers, cant say for certain which one is correct answer anymore. How can o
pass the exam like this????

upvoted 1 times

aliounegdiop 12 months ago

A. Organization Administrator

Here's why:

Organization Administrator: This role provides full control over all resources and policies within the organization, including permissions and auditing. It allows your team to manage permissions, policies, and configurations at the organizational level, making it the most appropriate choi when you need comprehensive control.

Security Reviewer: This role focuses on reviewing and assessing security configurations but doesn't grant the level of control needed for managin permissions and auditing at the organizational level.

Organization Role Administrator: This role allows management of IAM roles at the organization level but doesn't provide control over policies anc auditing.

Organization Policy Administrator: This role allows for the management of organization policies, but it doesn't cover permissions and auditing.

upvoted 3 times

elad17 1 year, 4 months ago

Selected Answer: A

A is the only role that gives you management permissions and not just viewing / role editing.

upvoted 4 times

Ishu_awsguy 1 year, 7 months ago

i would go with A. Audit of all domain resources might have a very broad scope and C might not have those permissions.
Because it is audit , i believe its a responsible job so A can be afforded

upvoted 2 times

GCP72 2 years ago

Selected Answer: C

The correct answer is C

upvoted 1 times

Medofree 2 years, 4 months ago

Answer is A, among the 4, it is the only role able de manage permissions

upvoted 3 times

Lancyqusa 2 years, 8 months ago

The answer must be A - check out the example that allows the CTO to setup permissions for the security team:
https://cloud.google.com/iam/docs/job-functions/auditing#scenario_operational_monitoring

upvoted 2 times

OSNG 3 years ago

Its A.
They are looking for Domain Resources Management i.e. Projects, Folders, Permissions. and only Organization Administrator is the only option allows it. Moreover, Organization Administrator is the only option that falls under "Used IN: Resource Manager"
roles/resourcemanager.organizationAdmin

upvoted 1 times

[Removed] 3 years, 5 months ago

C is the answer.
Here are the permissions available to organizationRoleAdmin

iam.roles.create
iam.roles.delete
iam.roles.undelete
iam.roles.get
iam.roles.list
iam.roles.update
resourcemanager.projects.get
resourcemanager.projects.getIamPolicy
resourcemanager.projects.list
resourcemanager.organizations.get
resourcemanager.organizations.getIamPolicy

There are sufficient as per least privilege policy. You can do user management as well as auditing.

upvoted 5 times

> **[Removed]** 3 years, 5 months ago
>
> link - https://cloud.google.com/iam/docs/understanding-custom-roles
>
> upvoted 1 times

DebasishLowes 3 years, 5 months ago

Ans : D. As it's related to Resources, so definitely policy comes into picture.

upvoted 1 times

HateMicrosoft 3 years, 6 months ago

Correct is D
https://cloud.google.com/resource-manager/docs/organization-policy/overview

upvoted 2 times

BhupalS 3 years, 8 months ago

Role Permissions
roles/iam.organizationRoleAdmin iam.roles.create
iam.roles.delete
iam.roles.undelete
iam.roles.get
iam.roles.list
iam.roles.update
resourcemanager.projects.get
resourcemanager.projects.getIamPolicy
resourcemanager.projects.list
resourcemanager.organizations.get
resourcemanager.organizations.getIamPolicy

upvoted 1 times

FatCharlie 3 years, 9 months ago

The confusion here, in my opinion, is that the question is asking for the ability to manage roles & audit _DOMAIN_ resources.

Domain resources in the GCP hierarchy are folders & projects, because those are the only things that can be directly under an Organization (aka Domain).

The Organization Role Admin is the option that gives you the ability to manage custom roles & list folders & projects.

upvoted 5 times

jonclem 3 years, 9 months ago

Organization Policy Administrator has 2 assigned permissions: orgpolicy.policy.get
orgpolicy.policy.set
Organization Role Administrator has 11 assigned permissions: iam.roles.create, iam.roles.delete, iam.roles.get, iam.roles.list, iam.roles.undelete, iam.roles.update, resourcemanager.organizations.get, resourcemanager.organizations.getIamPolicy, resourcemanager.projects.get, resourcemanager.projects.getIamPolicy, resourcemanager.projects.list,
While Security Review has over 700 permissions assigned to it.
With the question focusing on managing permissions and auditing I'd be inclined to go with option B: Security Reviewer.

upvoted 2 times

Question #15 *Topic 1*

An application running on a Compute Engine instance needs to read data from a Cloud Storage bucket. Your team does not allow Cloud Storage buckets to be globally readable and wants to ensure the principle of least privilege.
Which option meets the requirement of your team?

A. Create a Cloud Storage ACL that allows read-only access from the Compute Engine instance's IP address and allows the application to read from the bucket without credentials.

B. Use a service account with read-only access to the Cloud Storage bucket, and store the credentials to the service account in the config of the application on the Compute Engine instance.

C. Use a service account with read-only access to the Cloud Storage bucket to retrieve the credentials from the instance metadata.

D. Encrypt the data in the Cloud Storage bucket using Cloud KMS, and allow the application to decrypt the data with the KMS key.

Correct Answer: *C*

ESP_SAP Highly Voted 3 years, 9 months ago
Correct Answer is (B):
If your application runs inside a Google Cloud environment that has a default service account, your application can retrieve the service account credentials to call Google Cloud APIs. Such environments include Compute Engine, Google Kubernetes Engine, App Engine, Cloud Run, and Clouc Functions. We recommend using this strategy because it is more convenient and secure than manually passing credentials.

Additionally, we recommend you use Google Cloud Client Libraries for your application. Google Cloud Client Libraries use a library called Application Default Credentials (ADC) to automatically find your service account credentials. ADC looks for service account credentials in the following order:

https://cloud.google.com/docs/authentication/production#automatically
upvoted 13 times

ChewB666 3 years, 9 months ago
Hello guys!

Does anyone have the rest of the questions to share? :(
I can't see the rest of the issues because of the subscription.
upvoted 3 times

[Removed] 1 year, 1 month ago
Interestingly, the link you listed recommends using an attached service account. Attached service accounts use the metadata server to get credentials for the service.
Reference: https://cloud.google.com/docs/authentication/application-default-credentials#attached-sa
upvoted 2 times

[Removed] 1 year, 1 month ago
ADC tries to get credentials for attached service account from the environment variable first, then a "well-known location for credentials" (AKA Secret Manager) and then the metadata server. There is no reference for application configuration (i.e. code).
Which makes "B" invalid and "C" the correct choice.

https://cloud.google.com/docs/authentication/application-default-credentials#attached-sa
upvoted 2 times

Medofree Highly Voted 2 years, 4 months ago
Selected Answer: C
Correct ans is C. The credentials are retrieved from the metedata server
upvoted 13 times

okhascorpio Most Recent 10 months, 3 weeks ago
A. Although it would work, but it is less preferred method and are error prone.
B. Storing credentials in config is not good idea.
C. Is preferred method as applications can get credentials from instance metadata securely.
D. does not suggest controlled access, only encryption.
upvoted 2 times

ArizonaClassics 11 months, 3 weeks ago
C. Use a service account with read-only access to the Cloud Storage bucket to retrieve the credentials from the instance metadata.
upvoted 2 times

[Removed] 1 year, 1 month ago
Selected Answer: C

The answer is "C" because it references the preferred method for attaching a service account to an application.
The following page explains the preferred method for setting up a service account and attaching it to an application (where a metadata server is used to store credentials).
https://cloud.google.com/docs/authentication/application-default-credentials#attached-sa

upvoted 2 times

1br4in 1 year, 3 months ago

correct is B: Utilizzare un service account con accesso in sola lettura al bucket di Cloud Storage e archiviare le credenziali del service account nella configurazione dell'applicazione sull'istanza di Compute Engine.

Utilizzando un service account con accesso in sola lettura al bucket di Cloud Storage, puoi fornire all'applicazione le credenziali necessarie per leggere i dati dal bucket. Archiviando le credenziali del service account nella configurazione dell'applicazione sull'istanza di Compute Engine, garantisce che solo l'applicazione su quell'istanza abbia accesso alle credenziali e, di conseguenza, al bucket.

Questa opzione offre il principio del privilegio minimo, in quanto il service account ha solo i permessi necessari per leggere i dati dal bucket di Cloud Storage e le credenziali sono limitate all'applicazione specifica sull'istanza di Compute Engine. Inoltre, non richiede l'accesso globale ai bucket di Cloud Storage o l'utilizzo di autorizzazioni di accesso di rete basate su indirizzo IP.

upvoted 1 times

mahi9 1 year, 6 months ago

Selected Answer: C

C is the most viable option

upvoted 2 times

Meyucho 1 year, 9 months ago

Selected Answer: A

A CORRECT: It's the only answer when you use ACL to filter local IP's addresses and you can have the bucket without global access.
B INCORRET: Doesn't use the least privilege principle.
C INCORRECT: What credentials are we talking about!? To do this it's better option B.
D INCORRECT: Need global access.

upvoted 3 times

gcpengineer 1 year, 3 months ago

no.its not a soln

upvoted 1 times

dat987 1 year, 10 months ago

Selected Answer: B

meta data do not set service account

upvoted 2 times

[Removed] 1 year, 1 month ago

Application Default Credentials (ADC) is responsible for providing applications with credentials of the attached service account.
".. If ADC does not find credentials it can use in either the GOOGLE_APPLICATION_CREDENTIALS environment variable or the well-known location for Google Account credentials, it uses the metadata server to get credentials..."

https://cloud.google.com/docs/authentication/application-default-credentials#attached-sa

upvoted 2 times

GCP72 2 years ago

Selected Answer: C

The correct answer is C

upvoted 2 times

[Removed] 2 years, 5 months ago

B
If the environment variable GOOGLE_APPLICATION_CREDENTIALS is set, ADC uses the service account key or configuration file that the variable points to.
https://cloud.google.com/docs/authentication/production#automatically

upvoted 1 times

[Removed] 1 year, 1 month ago

"B" says "..config of the application.." which is stored in the code.
It does not say "environment variable".
Therefore the correct answer is "C" since credentials are also stored in metadata server too.

https://cloud.google.com/docs/authentication/application-default-credentials#attached-sa

upvoted 1 times

AaronLee 2 years, 5 months ago

The Answer is C
If the environment variable GOOGLE_APPLICATION_CREDENTIALS is set, ADC uses the service account key or configuration file that the variable points to.

If the environment variable GOOGLE_APPLICATION_CREDENTIALS isn't set, ADC uses the service account that is attached to the resource that is running your code.
https://cloud.google.com/docs/authentication/production#passing_the_path_to_the_service_account_key_in_code

upvoted 4 times

jj_618 2 years, 11 months ago

So is it B or C?

upvoted 1 times

StanPeng 2 years, 6 months ago

B for sure. C is wrong logic

upvoted 1 times

Medofree 2 years, 4 months ago

No the C is the right ans, you don't need to generate credentials into GCP since they are stored into metadata server, the application will retrieve them automatically through a Google Lib (or even manually by calling the url curl http://metadata.google.internal/computeMetadata/v1/instance/service-accounts/default/token -H "Metadata-Flavor: Google")

upvoted 3 times

Ishu_awsguy 1 year, 7 months ago

C is the right answer. If the service account has read permissions to cloud storage. Nothing extra is needed

upvoted 1 times

bolu 3 years, 7 months ago

Answer can be either B or C due to the relevance to servicing account. But storing password in app is a worst practice and we read it several times everywhere online hence it results in C as a best answer to handle service account through metadata

upvoted 5 times

[Removed] 1 year, 1 month ago

Agreed. B recommends storing credentials in code (app config) which is never good practice. Option C is the most secure out of all the option presented.

https://cloud.google.com/docs/authentication/application-default-credentials#attached-sa

upvoted 1 times

[Removed] 3 years, 10 months ago

Ans - C

upvoted 1 times

HectorLeon2099 3 years, 11 months ago

I'll go with B.
A - ACL's are not able to allow access based on IP
C - If you store the credentials in the metadata those will be public accessible by everyone with project access.
D - Too complex

upvoted 6 times

saurabh1805 3 years, 10 months ago

Yes B is best possible option. This is something google also recommnd.
https://cloud.google.com/storage/docs/authentication#libauth

upvoted 3 times

gcpengineer 1 year, 3 months ago

google never recommend that

upvoted 3 times

[Removed] 1 year, 1 month ago

B recommends storing credentials in code (app config) which is not recommended.
Correct answer is C.
Also metadata is different from metadata server. Metadata server is used to store service credentials for attached service accounts.

https://cloud.google.com/docs/authentication/application-default-credentials#attached-sa

upvoted 1 times

CHECK666 3 years, 11 months ago

c is correct

upvoted 2 times

Question #16 *Topic 1*

An organization's typical network and security review consists of analyzing application transit routes, request handling, and firewall rules. They want to enable their developer teams to deploy new applications without the overhead of this full review.
How should you advise this organization?

A. Use Forseti with Firewall filters to catch any unwanted configurations in production.

B. Mandate use of infrastructure as code and provide static analysis in the CI/CD pipelines to enforce policies.

C. Route all VPC traffic through customer-managed routers to detect malicious patterns in production.

D. All production applications will run on-premises. Allow developers free rein in GCP as their dev and QA platforms.

Correct Answer: *B*

bluetaurianbull Highly Voted 3 years, 5 months ago
@TNT87 and others, if you say (B) or even (C) or (A) can you provide proof and URLs to support your claims. Simply saying if you have done Clou
Architect you will know Everything under the sun is not the proper response, this is a discussion and a community here trying to learn. Not everyone will be in same standard or level.
Be helpful for others please....
upvoted 15 times

[Removed] 1 year, 1 month ago
Here you go for "B"
https://www.terraform.io/use-cases/enforce-policy-as-code
upvoted 1 times

OSNG Highly Voted 3 years ago
Its B.
Reasons:
1. They are asking for advise for Developers. (IaC is the suitable as they don't have to worry about managing infrastructure manually).
Moreover "An organization's typical network and security review consists of analyzing application transit routes, request handling, and firewall rules." statement is defining the process, they are not asking about the option to review the rules. Using Forseti is not reducing the overhead for Developers.
upvoted 10 times

ppandher Most Recent 10 months, 2 weeks ago
They want to enable their developer teams to deploy new applications without the overhead of this full review - Questions says this .
I am not sure if that feature is available in Forseti as per it, it is Inventory, Scanner, Explain, Enforce & Notification .
upvoted 1 times

[Removed] 1 year, 1 month ago
Selected Answer: B
The question emphasizes infrastructure related overhead. "B" is there only answer that addresses infrastructure overhead by leveraging infrastructure as code. Specifically the overhead is around security and policy concerns which are addressed by terraform in what they call "policy as code".

https://www.terraform.io/use-cases/enforce-policy-as-code
upvoted 1 times

TonytheTiger 1 year, 9 months ago
B: the best answer.
https://cloud.google.com/recommender/docs/tutorial-iac
upvoted 1 times

GCP72 2 years ago
Selected Answer: B
The correct answer is B
upvoted 1 times

Jeanphi72 2 years, 1 month ago
Selected Answer: A
The problem I see with B is that there is no reason why reviews should disappear: IaC is code and code needs to be reviewed before being deployed. Depending on the companies, devops writing terraform / CDK are named developers as well. Forseti seems to be able to automate this
https://github.com/forseti-security/forseti-security/tree/master/samples/scanner/scanners
upvoted 1 times

szl0144 2 years, 3 months ago
I think B is the answer, can anybody explain why A is correct?

upvoted 1 times

badrik 2 years, 3 months ago

A is detective in nature while B is preventive. So, It's B !

upvoted 2 times

minostrozaml2 2 years, 7 months ago

Took the tesk today, only 5 question from this dump, the rest are new questions.

upvoted 2 times

ThisisJohn 2 years, 8 months ago

Selected Answer: B

My vote goes to B by discard.

A) only mentions firewall rules, but nothing about network routes, and nothing on Forseti website either https://forsetisecurity.org/about/
C) Talks about malicious patterns, not about network routes, requests handling and patterns, like the question says
D) Running on-prem doesn't guarantee a higher level of control

Thus, the only answer that makes sense for me is B

upvoted 2 times

TNT87 3 years, 6 months ago

if you done Cloud Rchitect,you will understand why the answer is B

upvoted 4 times

bluetaurianbull 3 years, 4 months ago

its like saying if you have gone to space you experiance weighlessness .. be professional man... give proof for your claims, dont just expect world to be in same level as you. thats about COMMUNITY LEARNING ...

upvoted 10 times

TNT87 1 year, 5 months ago

kkkkkkkkkkkkk then research than being angry

upvoted 1 times

[Removed] 3 years, 10 months ago

Ans - C

upvoted 2 times

[Removed] 3 years, 10 months ago

Sry(Typo) .. It's B

upvoted 2 times

saurabh1805 3 years, 10 months ago

I will also go with option A

upvoted 1 times

CHECK666 3 years, 11 months ago

B is the answer

upvoted 1 times

ownez 4 years ago

Answer is B and not A because in A, the answer provided tells us the environment is in production where the question is about to enable their developer teams to deploy new applications without the overhead of the full review. Implementation of IAC is suitable for this.

Answer is B.

upvoted 3 times

MohitA 4 years ago

Yes B serves the purpose.

upvoted 2 times

aiwaai 4 years ago

Answer is A

upvoted 1 times

Question #17 *Topic 1*

An employer wants to track how bonus compensations have changed over time to identify employee outliers and correct earning disparities. This task must be performed without exposing the sensitive compensation data for any individual and must be reversible to identify the outlier.
Which Cloud Data Loss Prevention API technique should you use to accomplish this?

A. Generalization

B. Redaction

C. CryptoHashConfig

D. CryptoReplaceFfxFpeConfig

Correct Answer: *B*

xhova Highly Voted 👍 4 years, 5 months ago
Answer is D

https://cloud.google.com/dlp/docs/pseudonymization
upvoted 17 times

SilentSec 4 years, 1 month ago
Also the same usecase in the url that you post. D is right.
upvoted 1 times

smart123 4 years, 2 months ago
Option D is correct because it is reversible whereas option B is not.
upvoted 3 times

gcp_learner Highly Voted 👍 4 years, 2 months ago
The answer is A.
By bucketing or generalizing, we achieve a reversible pseudonymised data that can still yield the required analysis.
https://cloud.google.com/dlp/docs/concepts-bucketing
upvoted 6 times

Sheeda 4 years ago
Completely wrong

The answer is D for sure. The example was even in google docs but replaced for some reasons.

http://price2meet.com/gcp/docs/dlp_docs_pseudonymization.pdf
upvoted 7 times

crazycosmos Most Recent ⊙ 3 months, 1 week ago
Selected Answer: D
it is reversible for D
upvoted 1 times

ManuelY 3 months, 4 weeks ago
Selected Answer: D
Reversible
upvoted 1 times

Kiroo 4 months, 4 weeks ago
Selected Answer: D
For sure is D
https://cloud.google.com/sensitive-data-protection/docs/transformations-reference#fpe

I was in doubt about C but the hash can't be returned into the original value
upvoted 1 times

ketoza 7 months, 4 weeks ago
Selected Answer: D
https://cloud.google.com/dlp/docs/transformations-reference#fpe
upvoted 1 times

okhascorpio 10 months, 3 weeks ago
A. seems like good fit here. Preserve data utility while also reducing the identifiability of the data.
https://cloud.google.com/dlp/docs/concepts-bucketing

upvoted 1 times

okhascorpio 10 months, 3 weeks ago

I take it back. its not reversible.

upvoted 1 times

[Removed] 1 year, 1 month ago

Selected Answer: D

The keyword here is "reversible" or allows for "re-identification". Out of the options listed, Format preserving encryption (FPE-FFX) is the only one that allows "re-identification".

Therefore "D" is the most accurate option.

References:
https://cloud.google.com/dlp/docs/pseudonymization (see the table)
https://en.wikipedia.org/wiki/Format-preserving_encryption

upvoted 2 times

aashissh 1 year, 4 months ago

Selected Answer: A

Generalization is a technique that replaces an original value with a similar, but not identical, value. This technique can be used to help protect sensitive data while still allowing statistical analysis.

In this scenario, the employer can use generalization to replace the actual bonus compensation values with generalized values that are statisticall similar but not identical. This allows the employer to perform analysis on the data without exposing the sensitive compensation data for any individual employee.

Using Generalization can be reversible to identify outliers. The employer can then use the original data to investigate further and correct any earning disparities.

Redaction is another DLP API technique that can be used to protect sensitive data, but it is not suitable for this scenario since it would remove the data completely and make statistical analysis impossible. CryptoHashConfig and CryptoReplaceFfxFpeConfig are also not suitable for this scenaric since they are encryption techniques and do not allow statistical analysis of data.

upvoted 3 times

aashissh 1 year, 4 months ago

Answer is A:
Generalization is a technique that replaces an original value with a similar, but not identical, value. This technique can be used to help protect sensitive data while still allowing statistical analysis.

In this scenario, the employer can use generalization to replace the actual bonus compensation values with generalized values that are statisticall similar but not identical. This allows the employer to perform analysis on the data without exposing the sensitive compensation data for any individual employee.

Using Generalization can be reversible to identify outliers. The employer can then use the original data to investigate further and correct any earning disparities.

Redaction is another DLP API technique that can be used to protect sensitive data, but it is not suitable for this scenario since it would remove the data completely and make statistical analysis impossible. CryptoHashConfig and CryptoReplaceFfxFpeConfig are also not suitable for this scenaric since they are encryption techniques and do not allow statistical analysis of data.

upvoted 1 times

Lyfedge 1 year, 5 months ago

Correct Answer is (D): De-identifying sensitive data

Cloud Data Loss Prevention (DLP) can de-identify sensitive data in text content, including text stored in container structures such as tables. De-identification is the process of removing identifying information from data. The API detects sensitive data such as personally identifiable information (PII), and then uses a de-identification transformation to mask, delete, or otherwise obscure the data.

For example, de-identification techniques can include any of the following:

Masking sensitive data by partially or fully replacing characters with a symbol, such as an asterisk (*) or hash (#).

upvoted 1 times

mahi9 1 year, 6 months ago

Selected Answer: D

D is the most viable option

upvoted 1 times

null32sys 1 year, 6 months ago

The Answer is A

upvoted 1 times

Ishu_awsguy 1 year, 7 months ago

Correct answer is D. But,
The answer does not have a CryptoDeterministicConfig . We recommend using CryptoDeterministicConfig for all use cases which do not require preserving the input alphabet space and size, plus warrant referential integrity.

https://cloud.google.com/dlp/docs/transformations-reference#transformation_methods

upvoted 1 times

zanhsieh 1 year, 8 months ago

Answer D. Note that `CryptoReplaceFfxFpeConfig` might not be used in a real exam; they might change to `format preserve encryption`.

upvoted 5 times

Littleivy 1 year, 9 months ago

The answer is D
https://cloud.google.com/dlp/docs/transformations-reference#transformation_methods

upvoted 2 times

Premumar 1 year, 10 months ago

Selected Answer: D

D is the only option that is reversible.

upvoted 3 times

Question #18 *Topic 1*

An organization adopts Google Cloud Platform (GCP) for application hosting services and needs guidance on setting up password requirements for their Cloud
Identity account. The organization has a password policy requirement that corporate employee passwords must have a minimum number of characters.
Which Cloud Identity password guidelines can the organization use to inform their new requirements?

A. Set the minimum length for passwords to be 8 characters.

B. Set the minimum length for passwords to be 10 characters.

C. Set the minimum length for passwords to be 12 characters.

D. Set the minimum length for passwords to be 6 characters.

Correct Answer: *A*

bolu Highly Voted 3 years, 7 months ago
The situation changes year on year on GCP.Right now the right answer is C based on minimum requirement of 12 char in GCP as on Jan 2021.
https://support.google.com/accounts/answer/32040?hl=en#zippy=%2Cmake-your-password-longer-more-memorable
upvoted 15 times

desertlotus1211 3 years, 5 months ago
It asked for Cloud Indentity password requirements... Minimum is 8 Maximum is 100
upvoted 9 times

KILLMAD Highly Voted 4 years, 5 months ago
Ans is A
upvoted 12 times

rafaelc 4 years, 5 months ago
Default password length is 8 characters.
https://support.google.com/cloudidentity/answer/33319?hl=en
upvoted 10 times

pico Most Recent 3 months, 3 weeks ago
Selected Answer: D
Minimum is 6

https://cloud.google.com/identity-platform/docs/password-policy
upvoted 2 times

dija123 5 months ago
Selected Answer: A
Minimum 8
upvoted 1 times

madcloud32 5 months, 4 weeks ago
Selected Answer: C
12 is minimum good for app security.
upvoted 1 times

[Removed] 1 year, 1 month ago
Selected Answer: A
"A"
By default the minimum number of characters is 8 (max 100) however range can be adjusted.

https://support.google.com/a/answer/139399?sjid=18255262015630288726-NA
upvoted 2 times

amanshin 1 year, 2 months ago
Answer is A
The minimum password length for application hosting services on GCP was 12 characters until January 2023. However, it was recently changed to characters. This change was made to make it easier for users to create and remember strong passwords.
upvoted 1 times

Sachu555 1 year, 5 months ago
C is the correct ans

upvoted 1 times

Sammydp202020 1 year, 6 months ago

Selected Answer: A

Answer is A

upvoted 1 times

blue123456 1 year, 9 months ago

Ans A
https://support.google.com/cloudidentity/answer/2537800?hl=en#zippy=%2Creset-a-users-password

upvoted 2 times

xchmielu 1 year, 9 months ago

Selected Answer: C

https://support.google.com/accounts/answer/32040?hl=en#zippy=%2Cmake-your-password-longer-more-memorable

upvoted 1 times

GCP72 2 years ago

Selected Answer: A

The answer is A

upvoted 1 times

otokichi3 2 years, 2 months ago

The answer is A.
minimum character length is 8.
https://support.google.com/cloudidentity/answer/139399?hl=en

upvoted 1 times

Ksrp 2 years, 6 months ago

Ans - A, see https://support.google.com/accounts/answer/9094506#zippy=%2Cmake-your-password-longer-more-memorable

upvoted 1 times

umashankar_a 3 years, 1 month ago

Answer A
For Cloud Identity password requirements still is - Minimum 8 Maximum is 100
https://support.google.com/cloudidentity/answer/139399?hl=en#:~:text=It%20can%20be%20between%208,decide%20to%20change%20their%20password.

upvoted 3 times

soukumar369 3 years, 7 months ago

Answer is C : https://support.google.com/accounts/answer/32040?hl=en#zippy=%2Cmake-your-password-longer-more-memorable

Long passwords are stronger, so make your password at least 12 characters long. These tips can help you create longer passwords that are easier to remember. Try to use:

upvoted 2 times

> **DA95** 1 year, 8 months ago
>
> The question is about the minimum length, not about safer
>
> upvoted 1 times

[Removed] 3 years, 10 months ago

Ans - A

upvoted 3 times

Question #19 *Topic 1*

You need to follow Google-recommended practices to leverage envelope encryption and encrypt data at the application layer.
What should you do?

A. Generate a data encryption key (DEK) locally to encrypt the data, and generate a new key encryption key (KEK) in Cloud KMS to encrypt the DEK. Store both the encrypted data and the encrypted DEK.

B. Generate a data encryption key (DEK) locally to encrypt the data, and generate a new key encryption key (KEK) in Cloud KMS to encrypt the DEK. Store both the encrypted data and the KEK.

C. Generate a new data encryption key (DEK) in Cloud KMS to encrypt the data, and generate a key encryption key (KEK) locally to encrypt the key. Store both the encrypted data and the encrypted DEK.

D. Generate a new data encryption key (DEK) in Cloud KMS to encrypt the data, and generate a key encryption key (KEK) locally to encrypt the key. Store both the encrypted data and the KEK.

Correct Answer: *A*
Reference:
https://cloud.google.com/kms/docs/envelope-encryption

Sheeda Highly Voted 4 years ago
Yes, A is correct

The process of encrypting data is to generate a DEK locally, encrypt data with the DEK, use a KEK to wrap the DEK, and then store the encrypted data and the wrapped DEK. The KEK never leaves Cloud KMS.
upvoted 21 times

MohitA 4 years ago
Agree on A, spot on "KEK never leaves Cloud KMS"
upvoted 3 times

Di4sa Most Recent 6 months, 2 weeks ago
Selected Answer: A
A is the correct answer as stated in google docs
The process of encrypting data is to generate a DEK locally, encrypt data with the DEK, use a KEK to wrap the DEK, and then store the encrypted data and the wrapped DEK. The KEK never leaves Cloud KMS.
https://cloud.google.com/kms/docs/envelope-encryption#how_to_encrypt_data_using_envelope_encryption
upvoted 1 times

standm 1 year, 3 months ago
KMS is used for storing KEK in CSEK & CMEK
upvoted 1 times

aashissh 1 year, 4 months ago
Selected Answer: B
This follows the recommended practice of envelope encryption, where the DEK is encrypted with a KEK, which is managed by a KMS service such Cloud KMS. Storing both the encrypted data and the KEK allows for the data to be decrypted using the KEK when needed. It's important to generate the DEK locally to ensure the security of the key, and to generate a new KEK in Cloud KMS for added security and key management capabilities.
upvoted 1 times

ppandher 10 months, 2 weeks ago
We need to store the encrypted data and Wrapped DEK . KEK would be centrally Managed by KMS .
https://cloud.google.com/kms/docs/envelope-encryption#how_to_encrypt_data_using_envelope_encryption
upvoted 1 times

GCP72 2 years ago
Selected Answer: A
The answer is A
upvoted 2 times

minostrozaml2 2 years, 7 months ago
Took the tesk today, only 5 question from this dump, the rest are new questions.
upvoted 1 times

Bill831231 2 years, 8 months ago
A sounds like the correct answer:
https://cloud.google.com/kms/docs/envelope-encryption#how_to_encrypt_data_using_envelope_encryption

upvoted 1 times

umashankar_a 3 years, 1 month ago

Answer A
Envelope Encryption: https://cloud.google.com/kms/docs/envelope-encryption
Here are best practices for managing DEKs:
-Generate DEKs locally.
-When stored, always ensure DEKs are encrypted at rest.
- For easy access, store the DEK near the data that it encrypts.
The DEK is encrypted (also known as wrapped) by a key encryption key (KEK). The process of encrypting a key with another key is known as envelope encryption.
Here are best practices for managing KEKs:
-Store KEKs centrally. (KMS)
-Set the granularity of the DEKs they encrypt based on their use case. For example, consider a workload that requires multiple DEKs to encrypt th workload's data chunks. You could use a single KEK to wrap all DEKs that are responsible for that workload's encryption.
-Rotate keys regularly, and also after a suspected incident.

upvoted 2 times

desertlotus1211 3 years, 4 months ago

I'm no sure what the answers is, but the answers to this question has changed.... be prepared

upvoted 1 times

dtmtor 3 years, 5 months ago

Answer is A

upvoted 1 times

DebasishLowes 3 years, 5 months ago

Ans : A

upvoted 1 times

CloudTrip 3 years, 6 months ago

Correction I change it to A after reading the question once again.

upvoted 1 times

CloudTrip 3 years, 6 months ago

Answer is B as after DEK encryption it's KEK (not encrypted DEK) which never leaves KMS

upvoted 1 times

Bharathy 3 years, 9 months ago

A - Envelope Encryption (DEK - to encrypt the data, KEK - encrypt the DEK , KEK resides in KMS and only the encrypted data and wrapped DEK w be stored back)

upvoted 2 times

[Removed] 3 years, 10 months ago

Ans - A
https://cloud.google.com/kms/docs/envelope-encryption#how_to_encrypt_data_using_envelope_encryption

upvoted 1 times

CHECK666 3 years, 11 months ago

The answer is A

upvoted 1 times

aiwaai 4 years ago

The Answer is A

upvoted 2 times

Question #20 *Topic 1*

How should a customer reliably deliver Stackdriver logs from GCP to their on-premises SIEM system?

A. Send all logs to the SIEM system via an existing protocol such as syslog.

B. Configure every project to export all their logs to a common BigQuery DataSet, which will be queried by the SIEM system.

C. Configure Organizational Log Sinks to export logs to a Cloud Pub/Sub Topic, which will be sent to the SIEM via Dataflow.

D. Build a connector for the SIEM to query for all logs in real time from the GCP RESTful JSON APIs.

Correct Answer: *C*

ESP_SAP Highly Voted 3 years, 9 months ago

Correct answer is (C):
Scenarios for exporting Cloud Logging data: Splunk
This scenario shows how to export selected logs from Cloud Logging to Pub/Sub for ingestion into Splunk. Splunk is a security information and event management (SIEM) solution that supports several ways of ingesting data, such as receiving streaming data out of Google Cloud through Splunk HTTP Event Collector (HEC) or by fetching data from Google Cloud APIs through Splunk Add-on for Google Cloud.

Using the Pub/Sub to Splunk Dataflow template, you can natively forward logs and events from a Pub/Sub topic into Splunk HEC. If Splunk HEC is not available in your Splunk deployment, you can use the Add-on to collect the logs and events from the Pub/Sub topic.
https://cloud.google.com/solutions/exporting-stackdriver-logging-for-splunk

upvoted 18 times

AzureDP900 1 year, 10 months ago

I will go with C

upvoted 1 times

bkovari Most Recent 1 year ago

C is the only way to go

upvoted 2 times

GCP72 2 years ago

Selected Answer: C

I will go with C

upvoted 3 times

DebasishLowes 3 years, 5 months ago

Ans : C

upvoted 2 times

BlahBaller 3 years, 7 months ago

As I was the Logging Service Manager when we set this up with GCP. I can verify that C is how we have it setup, based on the Google's recommendations.

upvoted 2 times

Moss2011 3 years, 10 months ago

I think the correct one its D because C mention "Dataflow" and it cannot be connected to any sink out of GCP.

upvoted 1 times

[Removed] 3 years, 10 months ago

Ans - C
https://cloud.google.com/solutions/exporting-stackdriver-logging-for-splunk

upvoted 2 times

deevisrk 3 years, 10 months ago

C looks correct..
https://cloud.google.com/solutions/exporting-stackdriver-logging-for-splunk
Splunk is on premises SIEM solution in above example.

upvoted 2 times

saurabh1805 3 years, 10 months ago

I will go with Option B.

Read this email for more reason. C is not workable solution so that is first one not to consider.

upvoted 1 times

CHECK666 3 years, 11 months ago

C is the answer.

upvoted 1 times

ArizonaClassics 4 years, 1 month ago

I will go with C

upvoted 3 times

xhova 4 years, 5 months ago

C is correct

upvoted 4 times

Question #21 *Topic 1*

In order to meet PCI DSS requirements, a customer wants to ensure that all outbound traffic is authorized.
Which two cloud offerings meet this requirement without additional compensating controls? (Choose two.)

A. App Engine

B. Cloud Functions

C. Compute Engine

D. Google Kubernetes Engine

E. Cloud Storage

Correct Answer: *AC*
Reference:
https://cloud.google.com/solutions/pci-dss-compliance-in-gcp

KILLMAD Highly Voted 👍 4 years, 5 months ago
Answer is CD

because the doc mentions the following: "App Engine ingress firewall rules are available, but egress rules are not currently available:" and "Compute Engine and GKE are the preferred alternatives."
upvoted 17 times

rafaelc 4 years, 5 months ago
It is CD.
App Engine ingress firewall rules are available, but egress rules are not currently available. Per requirements 1.2.1 and 1.3.4, you must ensure that all outbound traffic is authorized. SAQ A-EP and SAQ D–type merchants must provide compensating controls or use a different Google Cloud product. Compute Engine and GKE are the preferred alternatives.

https://cloud.google.com/solutions/pci-dss-compliance-in-gcp
upvoted 6 times

Kiroo Most Recent 🕒 4 months, 4 weeks ago
Selected Answer: AB
Today this question does not have an specific answer it seems that compute engine and gke wound need additional steps to setup and functions and app engine it’s possible to just set egress so I would go with this pair
upvoted 2 times

techdsmart 6 months, 3 weeks ago
AB

With App Engine, you can ingress firewall rules and egress traffic controls .
You can use Cloud Functions ingress and egress network settings.
AB makes sense if we are talking about controlling ingress and egress traffic
upvoted 1 times

rottzy 11 months, 2 weeks ago
have a look 👀 at https://cloud.google.com/security/compliance/pci-dss#:~:text=The%20scope%20of%20the%20PCI,products%20against%20the%20PCI%20DSS.
there are multiple answers!
upvoted 1 times

GCBC 1 year ago
Ans is CD as per google docs - https://cloud.google.com/architecture/pci-dss-compliance-in-gcp#securing_your_network
upvoted 1 times

standm 1 year, 3 months ago
CD - since both support Egress firewalls.
upvoted 1 times

mahi9 1 year, 6 months ago
Selected Answer: CD
The most viable options
upvoted 1 times

civilizador 1 year, 6 months ago
Answer is CD:
https://cloud.google.com/architecture/pci-dss-compliance-in-gcp#securing_your_network

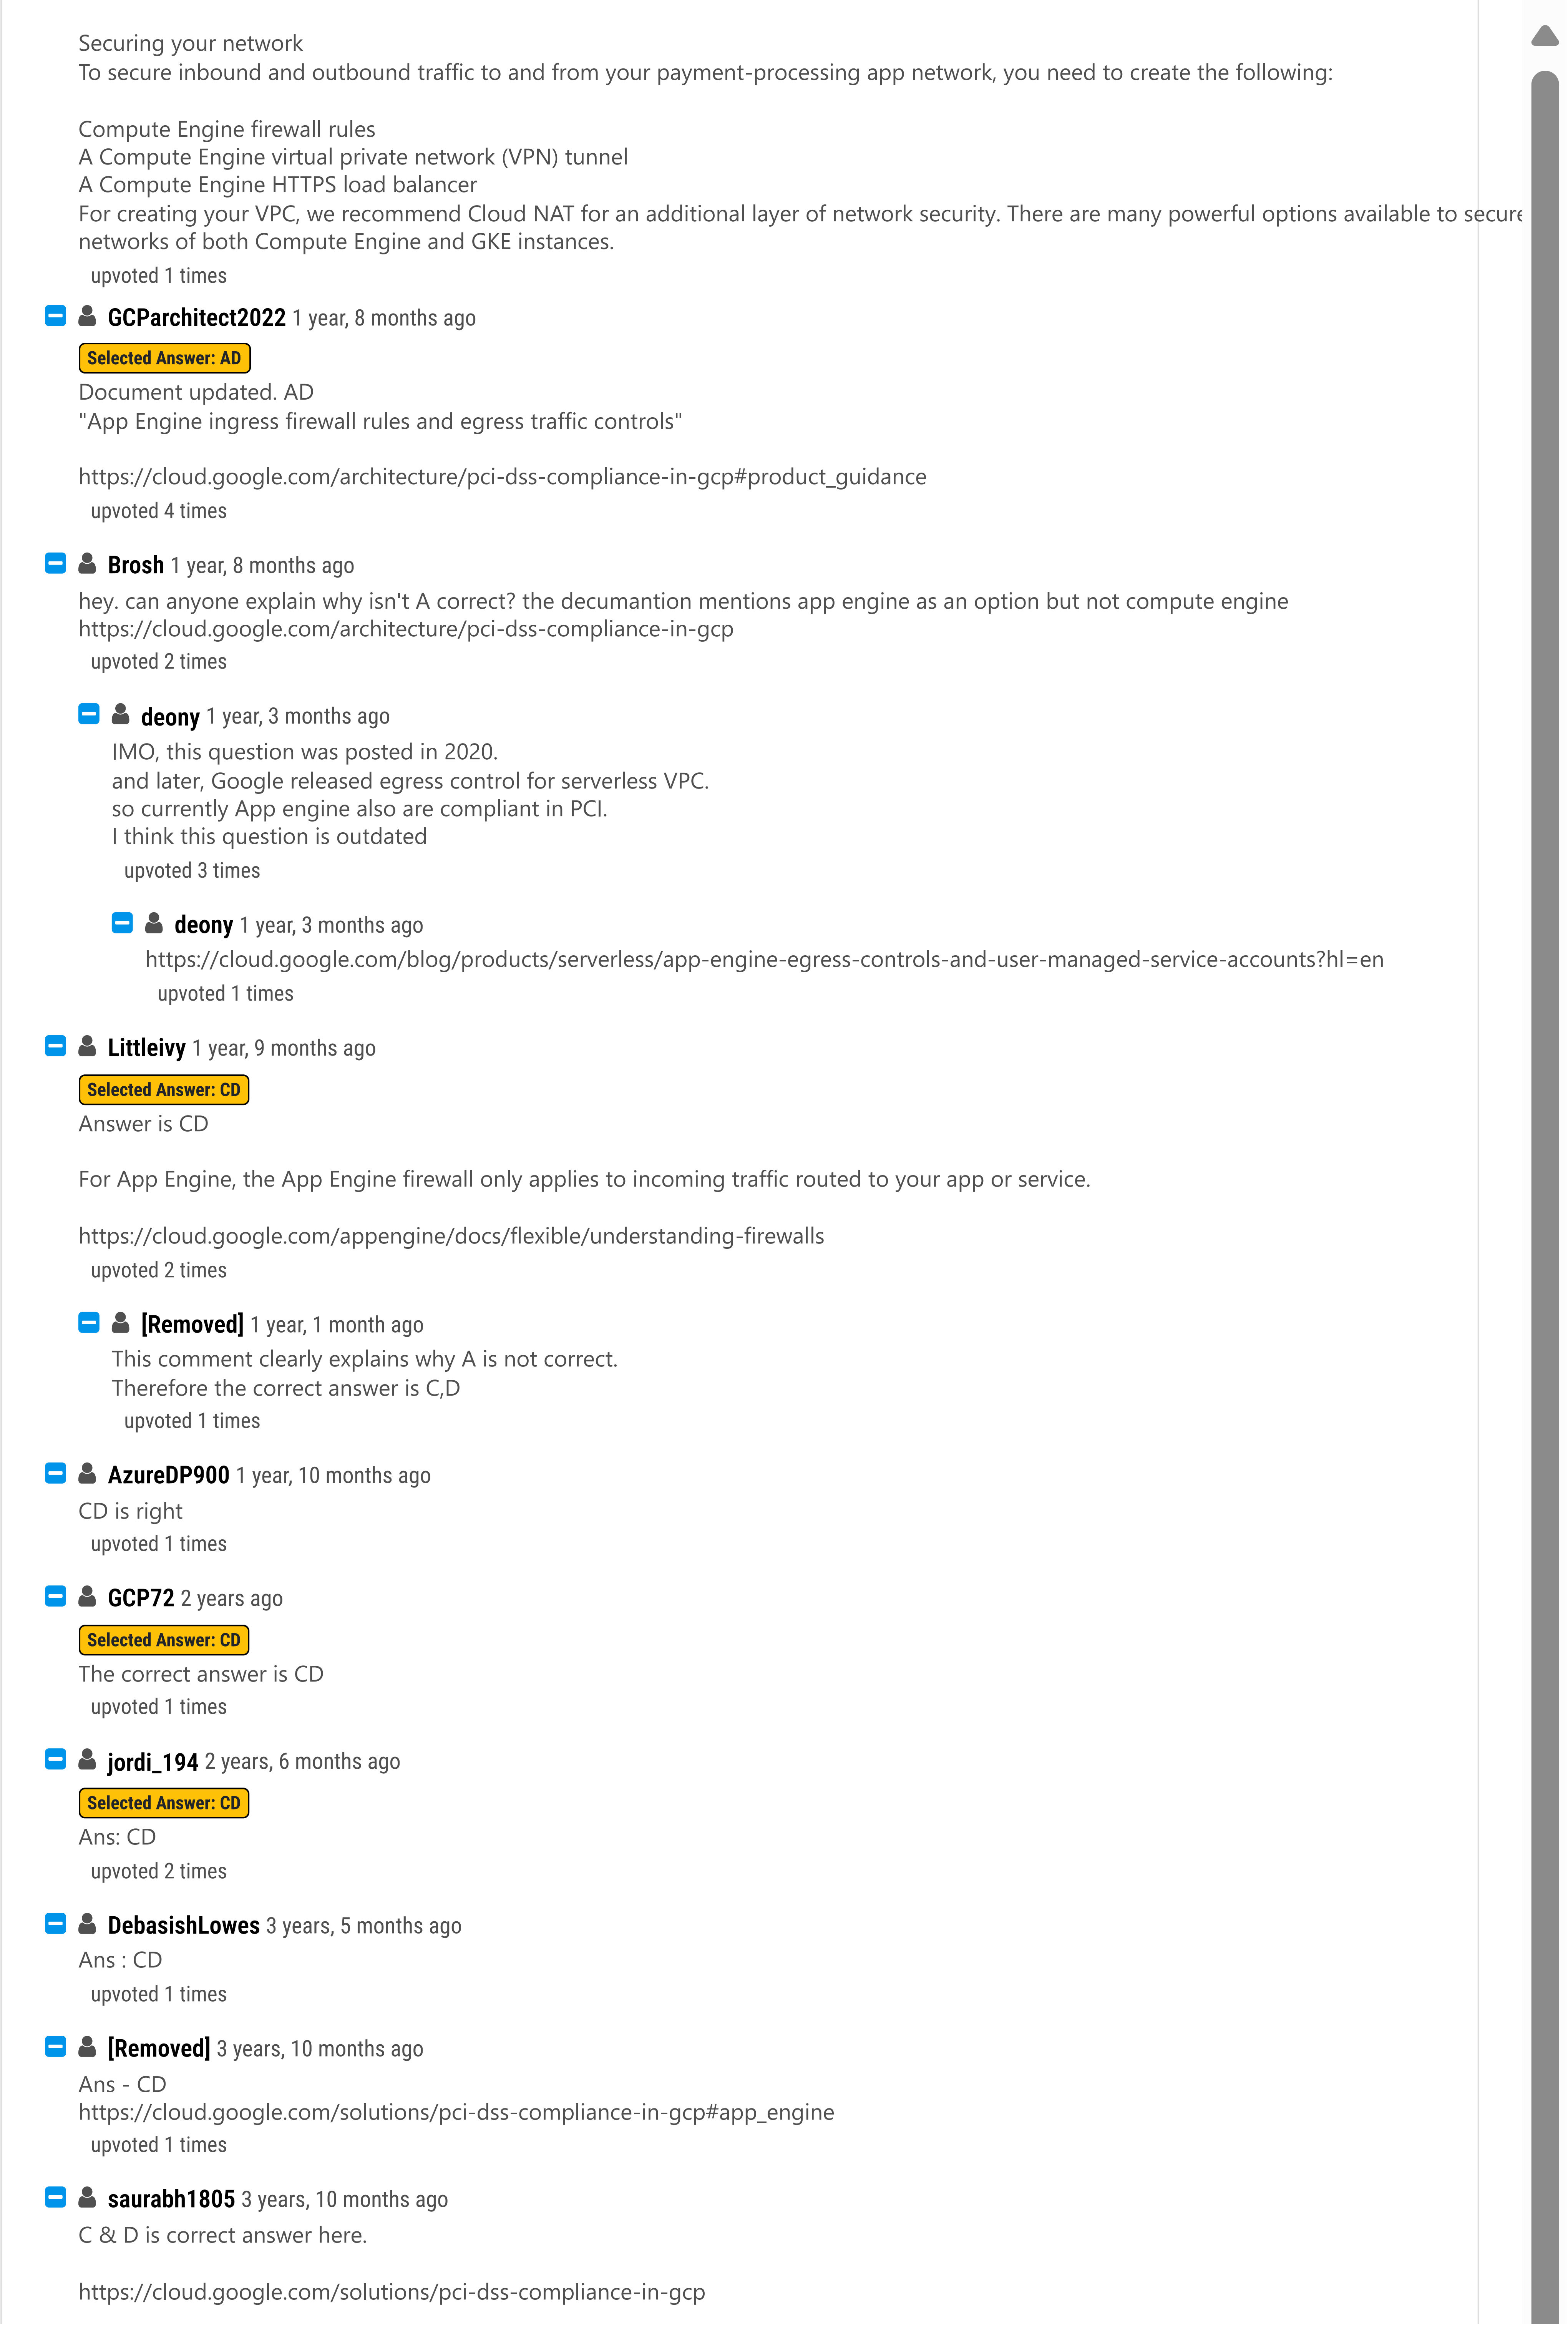

Securing your network
To secure inbound and outbound traffic to and from your payment-processing app network, you need to create the following:

Compute Engine firewall rules
A Compute Engine virtual private network (VPN) tunnel
A Compute Engine HTTPS load balancer
For creating your VPC, we recommend Cloud NAT for an additional layer of network security. There are many powerful options available to secure
networks of both Compute Engine and GKE instances.

upvoted 1 times

GCParchitect2022 1 year, 8 months ago

Selected Answer: AD

Document updated. AD
"App Engine ingress firewall rules and egress traffic controls"

https://cloud.google.com/architecture/pci-dss-compliance-in-gcp#product_guidance

upvoted 4 times

Brosh 1 year, 8 months ago

hey. can anyone explain why isn't A correct? the decumantion mentions app engine as an option but not compute engine
https://cloud.google.com/architecture/pci-dss-compliance-in-gcp

upvoted 2 times

deony 1 year, 3 months ago

IMO, this question was posted in 2020.
and later, Google released egress control for serverless VPC.
so currently App engine also are compliant in PCI.
I think this question is outdated

upvoted 3 times

deony 1 year, 3 months ago

https://cloud.google.com/blog/products/serverless/app-engine-egress-controls-and-user-managed-service-accounts?hl=en

upvoted 1 times

Littleivy 1 year, 9 months ago

Selected Answer: CD

Answer is CD

For App Engine, the App Engine firewall only applies to incoming traffic routed to your app or service.

https://cloud.google.com/appengine/docs/flexible/understanding-firewalls

upvoted 2 times

[Removed] 1 year, 1 month ago

This comment clearly explains why A is not correct.
Therefore the correct answer is C,D

upvoted 1 times

AzureDP900 1 year, 10 months ago

CD is right

upvoted 1 times

GCP72 2 years ago

Selected Answer: CD

The correct answer is CD

upvoted 1 times

jordi_194 2 years, 6 months ago

Selected Answer: CD

Ans: CD

upvoted 2 times

DebasishLowes 3 years, 5 months ago

Ans : CD

upvoted 1 times

[Removed] 3 years, 10 months ago

Ans - CD
https://cloud.google.com/solutions/pci-dss-compliance-in-gcp#app_engine

upvoted 1 times

saurabh1805 3 years, 10 months ago

C & D is correct answer here.

https://cloud.google.com/solutions/pci-dss-compliance-in-gcp

upvoted 1 times

Question #22 *Topic 1*

A website design company recently migrated all customer sites to App Engine. Some sites are still in progress and should only be visible to customers and company employees from any location.
Which solution will restrict access to the in-progress sites?

A. Upload an .htaccess file containing the customer and employee user accounts to App Engine.

B. Create an App Engine firewall rule that allows access from the customer and employee networks and denies all other traffic.

C. Enable Cloud Identity-Aware Proxy (IAP), and allow access to a Google Group that contains the customer and employee user accounts.

D. Use Cloud VPN to create a VPN connection between the relevant on-premises networks and the company's GCP Virtual Private Cloud (VPC) network.

Correct Answer: *C*

[Removed] Highly Voted 3 years, 10 months ago
Ans - C
https://cloud.google.com/iap/docs/concepts-overview#when_to_use_iap
upvoted 12 times

[Removed] Most Recent 1 year, 1 month ago
Selected Answer: C
This is the ideal use case for IAP.
"C" is the most accurate answer.
https://cloud.google.com/iap/docs/concepts-overview#when_to_use_iap
upvoted 1 times

GCP72 2 years ago
Selected Answer: C
The correct answer is C
upvoted 1 times

simbu1299 2 years, 5 months ago
Selected Answer: C
Answer is C
upvoted 1 times

mlx 3 years, 10 months ago
B - I think it is about to restrict access to 2 company networks, we can control access using IPs ranges, So Firewall rules should be sufficient. No need an extra product like IAP.. and also need users in Cloud Identity or other Idp federated..
upvoted 1 times

> **FatCharlie** 3 years, 9 months ago
> The sites should be accessible from any location, not just from the 2 company networks.
> upvoted 4 times

MohitA 4 years ago
C serves the purpose
upvoted 3 times

bigdo 4 years, 1 month ago
c is correct
upvoted 2 times

ArizonaClassics 4 years, 1 month ago
C is very correct
upvoted 2 times

SilentSec 4 years, 1 month ago
C is correct.
upvoted 2 times

Question #23 *Topic 1*

When working with agents in the support center via online chat, your organization's customers often share pictures of their documents with personally identifiable information (PII). Your leadership team is concerned that this PII is being stored as part of the regular chat logs, which are reviewed by internal or external analysts for customer service trends.
You want to resolve this concern while still maintaining data utility. What should you do?

A. Use Cloud Key Management Service to encrypt PII shared by customers before storing it for analysis.

B. Use Object Lifecycle Management to make sure that all chat records containing PII are discarded and not saved for analysis.

C. Use the image inspection and redaction actions of the DLP API to redact PII from the images before storing them for analysis.

D. Use the generalization and bucketing actions of the DLP API solution to redact PII from the texts before storing them for analysis.

Correct Answer: *C*
Reference:
https://cloud.google.com/dlp/docs/deidentify-sensitive-data

jitu028 Highly Voted 1 year, 11 months ago
Answer is C
upvoted 5 times

dija123 Most Recent 5 months, 2 weeks ago
Selected Answer: C
Agree with C
upvoted 1 times

standm 1 year, 3 months ago
since D talks about 'Text' and not image - it is not a suitable answer I guess.
upvoted 2 times

shayke 1 year, 8 months ago
Selected Answer: C
C the q refers to imaging
upvoted 4 times

kamal17 1 year, 8 months ago
Answer is C
upvoted 2 times

Question #24 *Topic 1*

A company's application is deployed with a user-managed Service Account key. You want to use Google-recommended practices to rotate the key.
What should you do?

A. Open Cloud Shell and run gcloud iam service-accounts enable-auto-rotate --iam-account=IAM_ACCOUNT.

B. Open Cloud Shell and run gcloud iam service-accounts keys rotate --iam-account=IAM_ACCOUNT --key=NEW_KEY.

C. Create a new key, and use the new key in the application. Delete the old key from the Service Account.

D. Create a new key, and use the new key in the application. Store the old key on the system as a backup key.

Correct Answer: *C*
Reference:
https://cloud.google.com/iam/docs/understanding-service-accounts

mdc Highly Voted 3 years, 2 months ago
C is correct. As explained, You can rotate a key by creating a new key, updating applications to use the new key, and deleting the old key. Use the serviceAccount.keys.create() method and serviceAccount.keys.delete() method together to automate the rotation.

https://cloud.google.com/iam/docs/creating-managing-service-account-keys#deleting_service_account_keys
upvoted 11 times

aliounegdiop Most Recent 12 months ago
B is correct. for C creating a new key and deleting the old one from the Service Account, is not recommended. Deleting the old key without replacing it could prevent your application from authenticating and accessing resources.
upvoted 1 times

aliounegdiop 12 months ago
my bad it should D. having a backup key in cae of problem with the new key
upvoted 1 times

eeghai7thioyaiR4 4 months, 1 week ago
If you keep the old key active, then your rotate is worthless (because anyone could still use the old key)

C is the solution: rotate and destroy the previous key
upvoted 3 times

[Removed] 1 year, 1 month ago
Selected Answer: C
"C" appears to be the most accurate.

https://cloud.google.com/iam/docs/key-rotation#process
upvoted 3 times

[Removed] 1 year, 1 month ago
"C" appears to be the most accurate.
https://cloud.google.com/iam/docs/key-rotation
upvoted 2 times

[Removed] 1 year, 1 month ago
Specifically: https://cloud.google.com/iam/docs/key-rotation#process
upvoted 1 times

megalucio 1 year, 2 months ago
Selected Answer: C
C it is the ans
upvoted 1 times

amanshin 1 year, 2 months ago
The correct answer is C. Create a new key, and use the new key in the application. Delete the old key from the Service Account.

Google recommends that you rotate user-managed service account keys every 90 days or less. This helps to reduce the risk of unauthorized acce
to your resources if the key is compromised.
upvoted 1 times

gcpengineer 1 year, 3 months ago
Selected Answer: C
C is the ans

upvoted 1 times

gcpengineer 1 year, 3 months ago

https://cloud.google.com/iam/docs/best-practices-for-managing-service-account-keys#rotate-keys

upvoted 1 times

aashissh 1 year, 4 months ago

Selected Answer: D

The recommended practice to rotate a user-managed Service Account key in GCP is to create a new key and use it in the application while keepir the old key for a specified period as a backup key. This helps to ensure that the application's service account always has a valid key and that there is no service disruption during the key rotation process. Therefore, the correct answer is option D.

upvoted 3 times

GCP72 2 years ago

Selected Answer: C

The correct answer is C

upvoted 2 times

absipat 2 years, 2 months ago

c of course

upvoted 1 times

DebasishLowes 3 years, 5 months ago

Ans : C

upvoted 2 times

[Removed] 3 years, 10 months ago

Ans - C
https://cloud.google.com/iam/docs/understanding-service-accounts#managing_service_account_keys

upvoted 4 times

ArizonaClassics 4 years ago

C is the right choice for me

upvoted 4 times

aiwaai 4 years ago

Correct Answer: C

upvoted 2 times

Question #25 *Topic 1*

Your team needs to configure their Google Cloud Platform (GCP) environment so they can centralize the control over networking resources like firewall rules, subnets, and routes. They also have an on-premises environment where resources need access back to the GCP resources through a private VPN connection.
The networking resources will need to be controlled by the network security team.
Which type of networking design should your team use to meet these requirements?

A. Shared VPC Network with a host project and service projects

B. Grant Compute Admin role to the networking team for each engineering project

C. VPC peering between all engineering projects using a hub and spoke model

D. Cloud VPN Gateway between all engineering projects using a hub and spoke model

Correct Answer: *A*
Reference:
https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations#centralize_network_control

ArizonaClassics Highly Voted 4 years, 1 month ago
I agree with A
Centralize network control:

Use Shared VPC to connect to a common VPC network. Resources in those projects can communicate with each other securely and efficiently across project boundaries using internal IPs. You can manage shared network resources, such as subnets, routes, and firewalls, from a central host project, enabling you to apply and enforce consistent network policies across the projects.
upvoted 19 times

ArizonaClassics 4 years, 1 month ago
WATCH: https://www.youtube.com/watch?v=WotV3D01tJA

READ:
https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations#centralize_network_control
upvoted 5 times

kamal17 Most Recent 1 year, 8 months ago
Answer is D , bocz On-prime user needs to access the GCP resources with help of Cloud VPN
upvoted 2 times

GCP72 2 years ago
Selected Answer: A
The correct answer is A
upvoted 1 times

minostrozaml2 2 years, 7 months ago
Took the tesk today, only 5 question from this dump, the rest are new questions.
upvoted 1 times

ZODOGAM 2 years, 9 months ago
Sheeda En mi caso te confirmo que desde la share VPC se establecen las VPNs y allí ingresa el tráfico desde los sitios locales. Definitivamente, la respuesta es la A
upvoted 1 times

DebasishLowes 3 years, 5 months ago
Ans : A. It will be shared VPC as it is asking for centralized network control.
upvoted 1 times

jonclem 3 years, 10 months ago
Option D is incorrect and a violation of Google's Service Specific terms as per : https://cloud.google.com/network-connectivity/docs/vpn/concepts/overview

I'd go with option A myself.
upvoted 1 times

[Removed] 3 years, 10 months ago
Ans - A
upvoted 1 times

saurabh1805 3 years, 10 months ago

A, this is exact reason to use shared VPC

upvoted 1 times

CHECK666 3 years, 11 months ago

A is the answer.

upvoted 1 times

Akku1614 4 years ago

A is correct as Shared VPC provides us with Centralized control however VPC Peering is a decentralized option.

upvoted 1 times

aiwaai 4 years ago

Correct Answer: A

upvoted 1 times

Sheeda 4 years ago

Connect your enterprise network

Many enterprises need to connect existing on-premises infrastructure with their Google Cloud resources. Evaluate your bandwidth, latency, and SLA requirements to choose the best connection option:

If you need low-latency, highly available, enterprise-grade connections that enable you to reliably transfer data between your on-premises and VPC networks without traversing the internet connections to Google Cloud, use Cloud Interconnect:

Dedicated Interconnect provides a direct physical connection between your on-premises network and Google's network.
Partner Interconnect provides connectivity between your on-premises and Google Cloud VPC networks through a supported service provider.
If you don't require the low latency and high availability of Cloud Interconnect, or you are just starting on your cloud journey, use Cloud VPN to s up encrypted IPsec VPN tunnels between your on-premises network and VPC. Compared to a direct, private connection, an IPsec VPN tunnel has lower overhead and costs.

upvoted 1 times

ArizonaClassics 4 years ago

Sheeda you need to read and understand the the question.

upvoted 1 times

ArizonaClassics 4 years ago

They are asking how you can centralize the control over networking resources like firewall rules, subnets, and routes. watch this:
https://www.youtube.com/watch?v=WotV3D01tJA
you will see that you can also manage vpn connections as well

upvoted 1 times

ESP_SAP 3 years, 9 months ago

you Should go back to the GCP Cloud Architect concepts or GCP Networking!

upvoted 2 times

Sheeda 4 years ago

I believe the answer is D. How can shared VPC give access to your on premise environment ? A seems wrong to me.

upvoted 4 times

AkbarM 1 year, 11 months ago

I also believe the same. i worked on interconnects and gateways to connect on prem resources.. only hub and spoke helps to connect onpremise network. ofcourse, we can centralize network controls using shared vpc. but the need here is some engineerng resources in on pre needs to access gcp resources. so this needs gateway to access gcp resources.

upvoted 2 times

Question #26 *Topic 1*

An organization is migrating from their current on-premises productivity software systems to G Suite. Some network security controls were in place that were mandated by a regulatory body in their region for their previous on-premises system. The organization's risk team wants to ensure that network security controls are maintained and effective in G Suite. A security architect supporting this migration has been asked to ensure that network security controls are in place as part of the new shared responsibility model between the organization and Google Cloud.
What solution would help meet the requirements?

A. Ensure that firewall rules are in place to meet the required controls.

B. Set up Cloud Armor to ensure that network security controls can be managed for G Suite.

C. Network security is a built-in solution and Google's Cloud responsibility for SaaS products like G Suite.

D. Set up an array of Virtual Private Cloud (VPC) networks to control network security as mandated by the relevant regulation.

Correct Answer: *C*

ESP_SAP Highly Voted 3 years, 9 months ago

Correct Answer is (C):
GSuite is Saas application.

Shared responsibility “Security of the Cloud” - GCP is responsible for protecting the infrastructure
that runs all of the services offered in the GCP Cloud. This infrastructure is composed of the hardware, software, networking, and facilities that rur
GCP Cloud services.

upvoted 11 times

AzureDP900 1 year, 10 months ago

C is right

upvoted 2 times

Topsy Highly Voted 3 years, 8 months ago

Answer is C- Review this Youtube Video- https://www.youtube.com/watch?v=D2zf0SgNdUw, scroll to 7:55, it would show you the Shared Responsibility model- With Gsuite being a SaaS product, Network Security is handled by Google

upvoted 7 times

okhascorpio Most Recent 6 months, 2 weeks ago

This thread suggests option "D" to be the only viable option. Now what ??
https://www.exam-answer.com/migrating-to-gsuite-network-security-controls

upvoted 1 times

[Removed] 1 year, 1 month ago

Selected Answer: C

GSuite AKA Workspace is software as a service where the SAAS provider (Google) is responsible for all underlying security.

https://youtu.be/D2zf0SgNdUw?t=535

upvoted 2 times

ppandey96 1 year, 5 months ago

Selected Answer: C

https://www.checkpoint.com/cyber-hub/cloud-security/what-is-google-cloud-platform-gcp-security/top-7-google-cloud-platform-gcp-security-best-practices/

upvoted 1 times

alleinallein 1 year, 5 months ago

Selected Answer: B

Shared responsibility model. Network security is not only Google's responsibility. As easy as that.

upvoted 1 times

Appsec977 1 year, 3 months ago

How would you set up a cloud armor in google workspace? totally misleading answer.

upvoted 3 times

alleinallein 1 year, 5 months ago

Need to change, as above if Google Workspace is considered as a Saas then network security is the responsibility of provider. C is correct.

upvoted 2 times

shayke 1 year, 8 months ago

Selected Answer: C

c - SAAS network security is the responsible of the cloud provider

upvoted 1 times

absipat 2 years, 2 months ago

c of course

upvoted 1 times

absipat 2 years, 2 months ago

C as it is SAAs

upvoted 1 times

FatCharlie 3 years, 9 months ago

Except for C, none of the options are possible in G Suite. There are no firewall, VPC, or Cloud Armor options there as far as I know.

upvoted 4 times

[Removed] 3 years, 10 months ago

Ans - A

upvoted 2 times

saurabh1805 3 years, 10 months ago

Question is asking for Network security group, Hence i will go with Option A

upvoted 1 times

skshak 3 years, 11 months ago

Answer is C. Gsuite is SaaS

upvoted 2 times

Question #27 *Topic 1*

A customer's company has multiple business units. Each business unit operates independently, and each has their own engineering group. Your team wants visibility into all projects created within the company and wants to organize their Google Cloud Platform (GCP) projects based on different business units. Each business unit also requires separate sets of IAM permissions.
Which strategy should you use to meet these needs?

A. Create an organization node, and assign folders for each business unit.

B. Establish standalone projects for each business unit, using gmail.com accounts.

C. Assign GCP resources in a project, with a label identifying which business unit owns the resource.

D. Assign GCP resources in a VPC for each business unit to separate network access.

Correct Answer: *A*

ArizonaClassics Highly Voted 4 years, 1 month ago
I will go with A
Refer to: https://cloud.google.com/resource-manager/docs/listing-all-resources

Also: https://wideops.com/mapping-your-organization-with-the-google-cloud-platform-resource-hierarchy/
upvoted 18 times

[Removed] Most Recent 1 year, 1 month ago
Selected Answer: A
"A"
Here's a blog post articulating this very business case.
https://cloud.google.com/blog/products/gcp/mapping-your-organization-with-the-google-cloud-platform-resource-hierarchy
upvoted 1 times

shayke 1 year, 8 months ago
Selected Answer: A
A is the right ans - resource manager
upvoted 1 times

DebasishLowes 3 years, 6 months ago
Ans - A
upvoted 3 times

[Removed] 3 years, 10 months ago
Ans - A
upvoted 1 times

aiwaai 4 years ago
Correct Answer: A
upvoted 1 times

Question #28 *Topic 1*

A company has redundant mail servers in different Google Cloud Platform regions and wants to route customers to the nearest mail server based on location.
How should the company accomplish this?

A. Configure TCP Proxy Load Balancing as a global load balancing service listening on port 995.

B. Create a Network Load Balancer to listen on TCP port 995 with a forwarding rule to forward traffic based on location.

C. Use Cross-Region Load Balancing with an HTTP(S) load balancer to route traffic to the nearest region.

D. Use Cloud CDN to route the mail traffic to the closest origin mail server based on client IP address.

Correct Answer: *D*

ESP_SAP Highly Voted 3 years, 9 months ago

Corrrect Answer is (A):

TCP Proxy Load Balancing is implemented on GFEs that are distributed globally. If you choose the Premium Tier of Network Service Tiers, a TCP proxy load balancer is global. In Premium Tier, you can deploy backends in multiple regions, and the load balancer automatically directs user traffi to the closest region that has capacity. If you choose the Standard Tier, a TCP proxy load balancer can only direct traffic among backends in a single region.

https://cloud.google.com/load-balancing/docs/load-balancing-overview#tcp-proxy-load-balancing

upvoted 26 times

Warren2020 Highly Voted 4 years, 1 month ago

A is the correct answer. D is not correct. CDN works with HTTP(s) traffic and requires caching, which is not a valid feature used for mail server

upvoted 9 times

Mr_MIXER007 Most Recent 6 days, 12 hours ago

Selected Answer: A

Corrrect Answer is (A)

upvoted 1 times

usercism007 2 weeks, 6 days ago

Select Answer: A

upvoted 1 times

3d9563b 1 month, 1 week ago

Selected Answer: A

TCP Proxy Load Balancing is the appropriate choice for globally routing TCP traffic, such as mail services, to the nearest server based on client location. It provides the necessary global load balancing capabilities to achieve this requirement.

upvoted 1 times

pico 3 months, 2 weeks ago

Selected Answer: B

why the other options are not the best fit:

A. TCP Proxy Load Balancing: This is a global load balancing solution, but it might not be the most efficient for routing mail traffic based on proximity.
C. Cross-Region Load Balancing with HTTP(S): This is designed for HTTP/HTTPS traffic, not mail protocols like POP3, SMTP, or IMAP.
D. Cloud CDN: While Cloud CDN can cache content for faster delivery, it's not designed to handle real-time mail traffic routing.

upvoted 1 times

shanwford 4 months, 1 week ago

Selected Answer: A

I go for (A) because Network Load Balancers are Layer 4 regional, passthrough load balancers: so it didnt work as global LB ("different GCP regions")

upvoted 1 times

eeghai7thioyaiR4 4 months, 1 week ago

This is probably an old question
2-3 years ago, GCP introduces a "proxy network load balancer"

So, in 2024, we have:
- application load balancer, global, external-only, multi-region backends, only for HTTP and HTTPS, do not preserve clients' IP
- "legacy" network load balancer (aka "passthrough"), external or internal, single-region, tcp or udp, preserve clients' IP
- "new" network load balancer (aka "proxy"), global, external or internal, multi-region backends, tcp or udp, do not preserve clients' IP

Here, we want:
- global
- external
- multi-region
- non-http
=> proxy network load balancer is the solution

This maps to A (generic answer) or B (but only in proxy mode: passthrough won't work)

upvoted 2 times

eeghai7thioyaiR4 4 months ago

On the other hand, B says "with forwarding rule". So this implies passthrough mode
This left only A as a solution

upvoted 1 times

Roro_Brother 4 months, 2 weeks ago

Selected Answer: B

The company can achieve location-based routing of customers to the nearest mail server in Google Cloud Platform (GCP) using a Network Load Balancer (NLB)

upvoted 1 times

dija123 6 months ago

Selected Answer: B

The company can achieve location-based routing of customers to the nearest mail server in Google Cloud Platform (GCP) using a Network Load Balancer (NLB)

upvoted 2 times

okhascorpio 6 months, 2 weeks ago

There is no direct SMTP support in TCP proxy load balancer, hens it cannot be A. Google Cloud best practices recommend Network Load Balancir (NLB) for Layer 4 protocols like SMTP.

upvoted 3 times

ErenYeager 6 months, 3 weeks ago

Selected Answer: B

B) Create a Network Load Balancer to listen on TCP port 995 with a forwarding rule to forward traffic based on location.

Explanation:

Port 995 implies this is SSL/TLS encrypted mail traffic (IMAP).
Network Load Balancing allows creating forwarding rules to route traffic based on IP location.
This can send users to the closest backend mail server.
TCP Proxy LB does not allow location-based routing.
HTTP(S) LB is for HTTP only, not generic TCP traffic.
Cloud CDN works at the HTTP level so cannot route TCP mail traffic.
So a Network Load Balancer with IP based forwarding rules provides the capability to direct mail users to the closest regional mail server based o their location, meeting the requirement.

upvoted 3 times

[Removed] 1 year, 1 month ago

Selected Answer: A

"A" is the most suitable answer.
Mail servers use SMTP which run on TCP. This excludes C, D which are HTTPs based. Option B is not global which excludes it as well.

The following page elaborates on global external proxy load balancing under the premium tier which meets the needs for this question and align with option A

https://cloud.google.com/load-balancing/docs/tcp#identify_the_mode

upvoted 4 times

gcpengineer 1 year, 3 months ago

Selected Answer: A

https://cloud.google.com/load-balancing/docs/tcp

upvoted 2 times

gcpengineer 1 year, 3 months ago

Selected Answer: B

B is the ans

upvoted 2 times

gcpengineer 1 year, 3 months ago

A is the ans. https://cloud.google.com/load-balancing/docs/tcp

upvoted 2 times

aashissh 1 year, 4 months ago

Selected Answer: B

The correct answer is B.

To route customers to the nearest mail server based on location, the company can create a Network Load Balancer. The Network Load Balancer ca listen on a specific TCP port (e.g., port 995 for mail traffic) and use a forwarding rule to forward traffic to the nearest mail server based on the client's location. This can be achieved by using a combination of the Load Balancing service and the Geo Map feature to route traffic based on th client's IP address.

TCP Proxy Load Balancing (A) is not suitable for this scenario as it is designed for non-HTTP(S) traffic, and it does not use client location information for traffic routing. Cross-Region Load Balancing (C) is also not suitable as it is designed for HTTP(S) traffic and does not use client location information for traffic routing. Cloud CDN (D) is designed for caching content and delivering it from the nearest point of presence (POP) to the user, but it does not route traffic to different servers based on the client's location.

upvoted 4 times

gcpengineer 1 year, 3 months ago

TCP proxy LB is relevant in this case

upvoted 2 times

mahi9 1 year, 6 months ago

Selected Answer: A

TCP Proxy Load Balancing is implemented on GFEs that are distributed globally. If you choose the Premium Tier of Network Service Tiers, a TCP proxy load balancer is global. In Premium Tier, you can deploy backends in multiple regions, and the load balancer automatically directs user traff to the closest region that has capacity. If you choose the Standard Tier, a TCP proxy load balancer can only direct traffic among backends in a single region.

upvoted 1 times

Question #29 *Topic 1*

Your team sets up a Shared VPC Network where project co-vpc-prod is the host project. Your team has configured the firewall rules, subnets, and VPN gateway on the host project. They need to enable Engineering Group A to attach a Compute Engine instance to only the 10.1.1.0/24 subnet.
What should your team grant to Engineering Group A to meet this requirement?

A. Compute Network User Role at the host project level.

B. Compute Network User Role at the subnet level.

C. Compute Shared VPC Admin Role at the host project level.

D. Compute Shared VPC Admin Role at the service project level.

Correct Answer: *C*
Reference:
https://cloud.google.com/vpc/docs/shared-vpc

mozammil89 Highly Voted 4 years, 5 months ago
The correct answer is B.

https://cloud.google.com/vpc/docs/shared-vpc#svc_proj_admins
upvoted 22 times

okhascorpio Most Recent 6 months, 2 weeks ago
Selected Answer: A
A is right. Source: https://cloud.google.com/compute/docs/access/iam#compute.networkUser
upvoted 1 times

stefanop 1 month, 3 weeks ago
this permission can be granted only at project level, not subnet level
upvoted 1 times

ErenYeager 6 months, 3 weeks ago
Selected Answer: B
B) Compute Network User Role at the subnet level.

The key points:

In a Shared VPC, the subnets are configured in the host project.
To allow another project to use a specific subnet, grant the Compute Network User role on that subnet.
The Compute Shared VPC Admin role allows full administration, which is more privileged than needed.
The Compute Network User role at the project level allows accessing all subnets, not just 10.1.1.0/24.
So granting the Compute Network User role specifically on the 10.1.1.0/24 subnet gives targeted access to only that subnet, meeting the requirement.
The subnet-level Compute Network User role provides the minimum necessary access to fulfill the need for Engineering Group A.
upvoted 4 times

Xoxoo 11 months, 2 weeks ago
Selected Answer: B
To enable Engineering Group A to attach a Compute Engine instance to only the 10.1.1.0/24 subnet in a Shared VPC setup, you should follow thes steps:

Grant the Compute Network User role at the service project level: This will allow members of Engineering Group A to create Compute Engine instances in their respective service projects.

Grant the Compute Network User role specifically on the 10.1.1.0/24 subnet: To ensure that Engineering Group A can only attach instances to the desired subnet, you should grant the Compute Network User role directly at the subnet level. This way, they have the necessary permissions for that specific subnet without impacting other subnets in the Shared VPC.

Option B, "Compute Network User Role at the subnet level," is the most appropriate choice in this scenario to achieve the desired outcome.
upvoted 3 times

shetniel 11 months, 2 weeks ago
The correct answer is B per least privilegd access rule
upvoted 2 times

[Removed] 1 year, 1 month ago
Selected Answer: B

"B" seems to be the most appropriate answer.
See step 4 here:
https://medium.com/google-cloud/google-cloud-shared-vpc-b33e0c9dd320

upvoted 2 times

aashissh 1 year, 4 months ago

Selected Answer: B

To enable Engineering Group A to attach a Compute Engine instance to only the 10.1.1.0/24 subnet in a Shared VPC Network where project co-vpc-prod is the host project, your team should grant Compute Network User Role at the subnet level. This will allow Engineering Group A to creat and manage resources in the specified subnet while restricting them from making changes to other resources in the host project. Granting Compute Network User Role at the host project level would allow Engineering Group A to create and manage resources across all subnets in the host project, which is more than what is needed in this case. Compute Shared VPC Admin Role at either the host or service project level would giv Engineering Group A too much control over the Shared VPC Network.

upvoted 2 times

mahi9 1 year, 6 months ago

Selected Answer: B

Admin role is not required

upvoted 2 times

Olen93 1 year, 6 months ago

The correct answer is B - https://cloud.google.com/compute/docs/access/iam#compute.networkUser states that the lowest level it can be grante on is project however I did confirm on my own companies shared VPC that roles/compute.networkUser can be granted at the subnet level

upvoted 1 times

amanp 1 year, 6 months ago

Selected Answer: A

Answer is A not B

The least level the Compute Network User role can be assigned is at Project level and NOT subnet level.

https://cloud.google.com/compute/docs/access/iam#compute.networkUser

upvoted 2 times

Meyucho 1 year, 9 months ago

Selected Answer: B

Grant network.user at subnet level:
https://cloud.google.com/vpc/docs/provisioning-shared-vpc#networkuseratsubnet

upvoted 2 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: B

The correct answer is B.

https://cloud.google.com/vpc/docs/shared-vpc#svc_proj_admins

upvoted 2 times

rajananna 1 year, 11 months ago

Selected Answer: A

Lowest level grant is at Project level. https://cloud.google.com/compute/docs/access/iam#compute.networkUser

upvoted 2 times

> **Premumar** 1 year, 10 months ago
>
> Lowest level grant is at Subnet level in this option. Project level is a broad level access.
>
> upvoted 2 times

tangac 1 year, 12 months ago

Selected Answer: A

based on that documentation it should clearly be done at the host project level :
https://cloud.google.com/compute/docs/access/iam#compute.networkUser

upvoted 3 times

piyush_1982 2 years, 1 month ago

Selected Answer: B

https://cloud.google.com/vpc/docs/shared-vpc#svc_proj_admins

upvoted 1 times

Medofree 2 years, 4 months ago

Selected Answer: B

The correct answer is b

upvoted 2 times

droppler 3 years, 1 month ago

The right one is b on my thinking, but i need to enable the other team to do the jobs, falls into D

upvoted 2 times

Question #30 *Topic 1*

A company migrated their entire data/center to Google Cloud Platform. It is running thousands of instances across multiple projects managed by different departments. You want to have a historical record of what was running in Google Cloud Platform at any point in time.
What should you do?

A. Use Resource Manager on the organization level.

B. Use Forseti Security to automate inventory snapshots.

C. Use Stackdriver to create a dashboard across all projects.

D. Use Security Command Center to view all assets across the organization.

Correct Answer: *B*

smart123 Highly Voted 👍 4 years, 2 months ago
'B is the correct answer. Only Forseti security can have both 'past' and 'present' (i.e. historical) records of the resources.
https://forsetisecurity.org/about/
upvoted 12 times

gcpengineer 1 year, 3 months ago
Forseti is outdated,no one uses it anymore
upvoted 4 times

mynk29 Highly Voted 👍 2 years, 6 months ago
Outdated questions- you should use asset inventory now.
upvoted 10 times

Roro_Brother Most Recent 🕓 4 months, 2 weeks ago
Selected Answer: D
D is good answer in this case. Foreseti is outdated
upvoted 2 times

Kiroo 4 months, 3 weeks ago
Selected Answer: D
It seems that for set is outdated and its features have been incorporated into security command center
upvoted 3 times

madcloud32 5 months, 4 weeks ago
Selected Answer: D
D is good answer in this case. Foreseti is outdated
upvoted 2 times

b6f53d8 8 months ago
D is a good answer
upvoted 2 times

ced3eals 10 months ago
Selected Answer: D
For an actual recent answer, D is the correct one.
upvoted 1 times

rottzy 11 months, 2 weeks ago
weird, Forseti - depreciated on Oct 2018, why was it even considered as an answer! 😉 😁
https://forsetisecurity.org/news/2019/02/18/deprecate-1.0.html
I'm going with option D
upvoted 1 times

cyberpunk21 1 year ago
Selected Answer: A
B is old way of doing things and things got updated
upvoted 2 times

[Removed] 1 year, 1 month ago
Selected Answer: B
"B" is the correct answer.
Forseti has been deprecated however it's capabilities and features (like asset inventory) have been incorporated into Security Command Center.

https://cloud.google.com/security-command-center/docs/concepts-security-command-center-overview#inventory

upvoted 1 times

amanshin 1 year, 2 months ago

Correct is A
Problem with Forseti - it's a third party tool, and it's sunset archived now due to lack of involvement. Do you really think Google would care to place it in test?

Using Resource Manager on the organization level is a good way to have a historical record of what was running in Google Cloud Platform at any point in time. This is because Resource Manager provides a centralized view of all of your organization's resources, including projects, folders, and organization policies. It's a native tool, so I would go for answer A.

upvoted 1 times

FunkyB 1 year, 7 months ago

B is the correct answer.
"Keep track of your environment
Take inventory snapshots of your Google Cloud Platform (GCP) resources on a recurring cadence so that you always have a history of what was in your cloud."

https://forsetisecurity.org/

upvoted 1 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: B

B is the correct answer. Only Forseti security can have both 'past' and 'present' (i.e. historical) records of the resources.
https://forsetisecurity.org/about/

upvoted 2 times

absipat 2 years, 2 months ago

b of course

upvoted 1 times

mitow95526 3 years, 3 months ago

https://cloud.google.com/security-command-center

Discover and view your assets in near-real time across App Engine, BigQuery, Cloud SQL, Cloud Storage, Compute Engine, Cloud Identity and Access Management, Google Kubernetes Engine, and more. Review historical discovery scans to identify new, modified, or deleted assets.

Why not D?

upvoted 4 times

> **ThisisJohn** 2 years, 8 months ago
>
> I guess the reason to discard D is that it says "all assets", while according to the documentation, "Security Command Center supports a large subset of Google Cloud assets.", so it supports a large number but not all assets.
>
> Ref: https://cloud.google.com/security-command-center/docs/concepts-security-command-center-overview#inventory
>
> upvoted 2 times
>
> **PATILDXB** 1 year, 8 months ago
>
> Azure security center does provide only realtime view on cloud. Endpoints once deleted or offboarded are no more visible in azure security center, which means historical details are lost
>
> upvoted 1 times

pfilourenco 3 years, 3 months ago

And about D?

upvoted 3 times

dtmtor 3 years, 5 months ago

Answer is B

upvoted 2 times

> **pythonrocks** 3 years, 1 month ago
>
> https://forsetisecurity.org/about/ inventory
>
> upvoted 1 times

Question #31 *Topic 1*

An organization is starting to move its infrastructure from its on-premises environment to Google Cloud Platform (GCP). The first step the organization wants to take is to migrate its current data backup and disaster recovery solutions to GCP for later analysis. The organization's production environment will remain on- premises for an indefinite time. The organization wants a scalable and cost-efficient solution.
Which GCP solution should the organization use?

A. BigQuery using a data pipeline job with continuous updates

B. Cloud Storage using a scheduled task and gsutil

C. Compute Engine Virtual Machines using Persistent Disk

D. Cloud Datastore using regularly scheduled batch upload jobs

Correct Answer: *A*

xhova Highly Voted 4 years, 5 months ago
Ans is B. A cost efficient disaster recovery solution is needed not a data warehouse.
upvoted 24 times

madcloud32 Most Recent 5 months, 4 weeks ago
Selected Answer: B
B is correct. It is about data backup, DR, not the database backup to GCP. BQ is not cost efficient compare to GCS
upvoted 1 times

tunstila 8 months ago
the two keywords here are 'later' and 'cost-efficient'. The company doesnt even know when the analysis will occur but they want to store the data
Storing it in BigQuery will not be cost-efficient for later analysis. Cloud Storage Archive is the best deal here.
upvoted 1 times

Nachtwaker 6 months ago
For later analysis means not now, so Bigquery is not required at this moment. Cloud storage content can be ingested in BigQuery 'later'. So should be B instead of A.
upvoted 1 times

W00kie 8 months, 3 weeks ago
Selected Answer: A
Imho A:
"The first step the organization wants to take is to migrate its current data backup and disaster recovery solutions to GCP for later analysis"
both solutions are scalable and cost efficient, but cloud storage is not designed for queirng, therefore data analysis would be easier in BigQuery.
upvoted 1 times

[Removed] 1 year, 1 month ago
Selected Answer: B
The keyword in the question here is "cost-effective".
Out of the 3 Disaster Recovery patterns (Cold, Warm, Hot HA), Cold is the most cost-effective which utilizes cloud storage.

References:
https://cloud.google.com/architecture/dr-scenarios-for-applications#cold-pattern-recovery-to-gcp

https://cloud.google.com/architecture/dr-scenarios-planning-guide#use-cloud-storage-as-part-of-your-daily-backup-routine
upvoted 2 times

raj117 1 year, 1 month ago
Right Answer is B
upvoted 2 times

SMB2022 1 year, 1 month ago
Correct Answer: B
upvoted 2 times

AwesomeGCP 1 year, 11 months ago
Selected Answer: B
B confirmed :-) https://cloud.google.com/solutions/dr-scenarios-planning-guide#use-cloud-storage-as-part-of-your-daily-backup-routine
upvoted 3 times

AzureDP900 1 year, 10 months ago
It is B

upvoted 2 times

giovy_82 2 years ago

I would go for B, but a doubt remains: it is talking about Disaster Recovery solution, which could not only be related to data but also to VM and applications running inside VMs. any way B is more cost-efficient than A, considering also that data backup need to be moved to GCP.

upvoted 1 times

absipat 2 years, 2 months ago

B of course

upvoted 2 times

DebasishLowes 3 years, 5 months ago

Ans : B. Cloud storage is cost efficient one.

upvoted 4 times

[Removed] 3 years, 10 months ago

Ans - B

upvoted 2 times

CHECK666 3 years, 11 months ago

B is the answer.

upvoted 2 times

paxjoshi 4 years ago

B is the correct answer. They need the data for later analysis and they are looking for cost-effective service.

upvoted 2 times

aiwaai 4 years ago

Correct Answer: A

upvoted 1 times

> **aiwaai** 4 years ago
>
> I make corrections, B is Correct Answer.
>
> upvoted 1 times

ArizonaClassics 4 years, 1 month ago

Answer B works for me as the type of workload to be stored is not stated or defined

upvoted 1 times

SilentSec 4 years, 1 month ago

B confirmed: https://cloud.google.com/solutions/dr-scenarios-planning-guide#use-cloud-storage-as-part-of-your-daily-backup-routine

upvoted 3 times

Question #32 *Topic 1*

You are creating an internal App Engine application that needs to access a user's Google Drive on the user's behalf. Your company does not want to rely on the current user's credentials. It also wants to follow Google-recommended practices.
What should you do?

A. Create a new Service account, and give all application users the role of Service Account User.

B. Create a new Service account, and add all application users to a Google Group. Give this group the role of Service Account User.

C. Use a dedicated G Suite Admin account, and authenticate the application's operations with these G Suite credentials.

D. Create a new service account, and grant it G Suite domain-wide delegation. Have the application use it to impersonate the user.

Correct Answer: *A*

mozammil89 Highly Voted 👍 4 years, 5 months ago
I think the correct answer is D

https://developers.google.com/admin-sdk/directory/v1/guides/delegation
upvoted 16 times

eeghai7thioyaiR4 Most Recent 4 months, 1 week ago
A and B are wrong
Service Account User is use to grant someone the ability to impersonate a service account (ref: https://cloud.google.com/iam/docs/understandin
roles)

So with those solution, the user could do some actions as the newly created service account
We want the opposite: the service account need to do some actions as some user
=> D is the only working solution
upvoted 1 times

chagchoug 6 months, 3 weeks ago
Selected Answer: D
Option A is false because it does not address the requirement of accessing a user's Google Drive on their behalf without relying on the user's credentials. Instead, option D, which involves granting domain-wide delegation to a service account for impersonation, is the recommended approach for this scenario.
upvoted 1 times

Olen93 1 year, 6 months ago
I'm not sure if D is the correct answer. The question specifically states that they want to follow Google-recommended practices and https://cloud.google.com/iam/docs/best-practices-service-accounts#domain-wide-delegation states to avoid domain-wide delegation. I do agre
that D is the only way a service account can impersonate the user though
upvoted 1 times

Meyucho 1 year, 9 months ago
Selected Answer: D
A (Wrong) The access will be with the SA not the user's account.
B (Wrong) Same as A.
C. (Wrong) In this case the access is with the admins account, not user's.
D. (CORRECT!) It's the only answer that really impersonate the user.
upvoted 3 times

AzureDP900 1 year, 10 months ago
D. Create a new service account, and grant it G Suite domain-wide delegation. Have the application use it to impersonate the user.
upvoted 1 times

AwesomeGCP 1 year, 11 months ago
Selected Answer: D
correct answer is D
https://developers.google.com/admin-sdk/directory/v1/guides/delegation
upvoted 2 times

Medofree 2 years, 4 months ago
Selected Answer: D
Clearly D is the right answer
upvoted 2 times

Rhehehe 2 years, 8 months ago
They are asking for google recommended practice. Does D says that?

upvoted 1 times

[Removed] 3 years, 10 months ago

Ans - D

upvoted 2 times

CHECK666 3 years, 11 months ago

D is the answer.

upvoted 1 times

ArizonaClassics 4 years, 1 month ago

D is the best choice

upvoted 1 times

MarkDillon1075 4 years, 2 months ago

I agree D

upvoted 1 times

Question #33

Topic 1

A customer wants to move their sensitive workloads to a Compute Engine-based cluster using Managed Instance Groups (MIGs). The jobs are bursty and must be completed quickly. They have a requirement to be able to control the key lifecycle.
Which boot disk encryption solution should you use on the cluster to meet this customer's requirements?

A. Customer-supplied encryption keys (CSEK)

B. Customer-managed encryption keys (CMEK) using Cloud Key Management Service (KMS)

C. Encryption by default

D. Pre-encrypting files before transferring to Google Cloud Platform (GCP) for analysis

Correct Answer: *B*

Reference -

https://cloud.google.com/kubernetes-engine/docs/how-to/dynamic-provisioning-cmek

animesh54 Highly Voted 👍 2 years, 4 months ago

Selected Answer: B

Customer Managed Encryption keys using KMS lets users control the key management and rotation policies and Compute Engine Disks support CMEKs

upvoted 6 times

AwesomeGCP Highly Voted 👍 1 year, 11 months ago

Selected Answer: B

Correct Answer: B
Explanation/Reference:
Reference https://cloud.google.com/kubernetes-engine/docs/how-to/dynamic-provisioning-cmek

upvoted 5 times

trashbox Most Recent 🕓 4 months ago

Selected Answer: B

"Control over the key lifecycle" is the key. The KMS is the most appropriate solution.

upvoted 1 times

Question #34 *Topic 1*

Your company is using Cloud Dataproc for its Spark and Hadoop jobs. You want to be able to create, rotate, and destroy symmetric encryption keys used for the persistent disks used by Cloud Dataproc. Keys can be stored in the cloud.
What should you do?

A. Use the Cloud Key Management Service to manage the data encryption key (DEK).

B. Use the Cloud Key Management Service to manage the key encryption key (KEK).

C. Use customer-supplied encryption keys to manage the data encryption key (DEK).

D. Use customer-supplied encryption keys to manage the key encryption key (KEK).

Correct Answer: *A*

mte_tech34 Highly Voted 👍 3 years, 11 months ago
Answer is B.
https://cloud.google.com/dataproc/docs/concepts/configuring-clusters/customer-managed-encryption
"The CMEK feature allows you to create, use, and revoke the key encryption key (KEK). Google still controls the data encryption key (DEK)."
upvoted 25 times

passtest100 3 years, 11 months ago
SHOULD BE A.
NO envelope encryption is metioned in the question.
upvoted 5 times

Arad 2 years, 9 months ago
Correct answer is B, and A is wrong!
envlope encryption is default mechanism in CMEK when used for Dataproc, please check this link:

This PD and bucket data is encrypted using a Google-generated data encryption key (DEK) and key encryption key (KEK). The CMEK featur
allows you to create, use, and revoke the key encryption key (KEK). Google still controls the data encryption key (DEK). For more informatio
on Google data encryption keys, see Encryption at Rest.

https://cloud.google.com/dataproc/docs/concepts/configuring-clusters/customer-managed-encryption
upvoted 2 times

mynk29 2 years, 6 months ago
I agree but then should answer not be be C- customer supplied key?
upvoted 1 times

mynk29 2 years, 6 months ago
My bad I read it as Customer managed.. even though i now realised i wrote customer supplied. :D
upvoted 1 times

Sarmee305 Most Recent 2 months, 3 weeks ago
Selected Answer: B
Answer is B
Cloud KMS allows you to manage KEKs, which in turn are used to encrypt the DEKs. DEKs are then used to encrypt the data. This separation ensures that the more sensitive KEK remains securely managed within the Cloud KMS
upvoted 1 times

dija123 5 months ago
Selected Answer: B
Agree with B
upvoted 1 times

amanshin 1 year, 2 months ago
The correct answer is B. Use the Cloud Key Management Service to manage the key encryption key (KEK).

Cloud Dataproc uses a two-level encryption model, where the data encryption key (DEK) is encrypted with a key encryption key (KEK). The KEK is stored in Cloud Key Management Service (KMS), which allows you to create, rotate, and destroy the KEK as needed.

If you use customer-supplied encryption keys (CSEKs) to manage the DEK, you will be responsible for managing the CSEKs yourself. This can be a complex and time-consuming task, and it can also increase the risk of data loss if the CSEKs are compromised.
upvoted 1 times

aashissh 1 year, 4 months ago
Selected Answer: A

Option B, using Cloud KMS to manage the key encryption key (KEK), is not necessary as persistent disks in Cloud Dataproc are already encrypted rest using AES-256 encryption with a unique DEK generated and managed by Google.

upvoted 1 times

mahi9 1 year, 6 months ago

Selected Answer: B

The CMEK feature allows you to create, use, and revoke the key encryption key (KEK). Google still controls the data encryption key (DEK)."

upvoted 1 times

sameer2803 1 year, 6 months ago

there is a diagram in the link. if you understand the diagram, you will get the answer. https://cloud.google.com/sql/docs/mysql/cmek#with-cmek

upvoted 1 times

sameer2803 1 year, 6 months ago

Answer is B. the documentation says that Google does the data encryption by default and then that encryption key is again encrypted by KEK. which in turn can be managed by Customer.

upvoted 1 times

DA95 1 year, 8 months ago

Selected Answer: A

Option B, using the Cloud KMS to manage the key encryption key (KEK), is incorrect. The KEK is used to encrypt the DEK, so the DEK is the key tha is managed by the Cloud KMS.

upvoted 1 times

Meyucho 1 year, 9 months ago

Selected Answer: A

B can be right but we never been asked about envelope encription... so... the solution is to use a customer managed Data Encryption Key

upvoted 1 times

AzureDP900 1 year, 10 months ago

B. Use the Cloud Key Management Service to manage the key encryption key (KEK).

upvoted 1 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: B

Answer is B,

https://cloud.google.com/dataproc/docs/concepts/configuring-clusters/customer-managed-encryption

upvoted 4 times

giovy_82 2 years ago

Selected Answer: B

In my opinion it should be B. reference :
https://cloud.google.com/kms/docs/envelope-encryption
How to encrypt data using envelope encryption
The process of encrypting data is to generate a DEK locally, encrypt data with the DEK, use a KEK to wrap the DEK, and then store the encrypted data and the wrapped DEK. The KEK never leaves Cloud KMS.

upvoted 2 times

piyush_1982 2 years, 1 month ago

Selected Answer: A

I think the answer is A.
DEK (Data encryption Key) is the key which is used to encrypt the data. It can be both customer-managed or customer supplied in terms of GCP>
https://cloud.google.com/dataproc/docs/concepts/configuring-clusters/customer-managed-encryption

The link above states "This PD and bucket data is encrypted using a Google-generated data encryption key (DEK) and key encryption key (KEK). T CMEK feature allows you to create, use, and revoke the key encryption key (KEK). Google still controls the data encryption key (DEK)."

upvoted 1 times

absipat 2 years, 2 months ago

b of course

upvoted 1 times

[Removed] 3 years, 4 months ago

I also support B, but A is also good ,because kek is hosted within KMS, also the real DEK can be uploaded there ,or just in the database.

upvoted 2 times

DebasishLowes 3 years, 5 months ago

Ans : B.

upvoted 2 times

Question #35 *Topic 1*

You are a member of the security team at an organization. Your team has a single GCP project with credit card payment processing systems alongside web applications and data processing systems. You want to reduce the scope of systems subject to PCI audit standards.
What should you do?

A. Use multi-factor authentication for admin access to the web application.

B. Use only applications certified compliant with PA-DSS.

C. Move the cardholder data environment into a separate GCP project.

D. Use VPN for all connections between your office and cloud environments.

Correct Answer: *C*

jonclem Highly Voted 4 years, 5 months ago
I'd go for answer C myself.

https://cloud.google.com/solutions/best-practices-vpc-design
upvoted 22 times

[Removed] Highly Voted 3 years, 10 months ago
Ans - C
https://cloud.google.com/solutions/pci-dss-compliance-in-gcp#setting_up_your_payment-processing_environment
upvoted 7 times

AzureDP900 Most Recent 1 year, 10 months ago
answer is C
upvoted 1 times

Medofree 2 years, 4 months ago
Selected Answer: C
Projets are units of isolationm the answer is C.
upvoted 2 times

CHECK666 3 years, 11 months ago
C is the answer.
upvoted 1 times

smart123 4 years, 2 months ago
The Answer is C. Check "Setting up your payment-processing environment" section in https://cloud.google.com/solutions/pci-dss-compliance-in-gcp.
In the question, it is mentioned that it is the same environment for card processing as the Web App and Data processing and that is not recommended.
upvoted 4 times

xhova 4 years, 5 months ago
Definitely C
upvoted 1 times

Question #36 *Topic 1*

A retail customer allows users to upload comments and product reviews. The customer needs to make sure the text does not include sensitive data before the comments or reviews are published.
Which Google Cloud Service should be used to achieve this?

A. Cloud Key Management Service

B. Cloud Data Loss Prevention API

C. BigQuery

D. Web Security Scanner

Correct Answer: *D*

rafaelc Highly Voted 4 years, 5 months ago
It's definitely B. It was on the practice test on google site.
B. Cloud Data Loss Prevention API
upvoted 28 times

Sarmee305 Most Recent 2 months, 3 weeks ago
B. Cloud Data Loss Prevention API
upvoted 1 times

uiuiui 10 months ago
Selected Answer: B
correct is B
upvoted 1 times

alleinallein 1 year, 5 months ago
Selected Answer: B
DLP is the only reasonable answer here. Security Scan is connected to AppSec.
upvoted 1 times

VishalBulbule 1 year, 8 months ago
"before the comments or reviews are published" - how will we use DLP API , so web scanner can be considered for correct answer.
upvoted 1 times

huntergame 1 year, 10 months ago
Selected Answer: B
Its obvious DLP
upvoted 1 times

PopeyeTheSailorMan 2 years, 1 month ago
Selected Answer: B
The answer can not be D (I am laughing loud since I use D for the reason of security scanning) hence the correct answer is B and it is not D
upvoted 1 times

Bwitch 2 years, 9 months ago
Selected Answer: B
DLP provides the service of redaction.
upvoted 3 times

DebasishLowes 3 years, 6 months ago
Its B.
upvoted 2 times

saurabh1805 3 years, 10 months ago
B is correct answer here.
upvoted 2 times

[Removed] 3 years, 10 months ago
Ans - B
upvoted 2 times

CHECK666 3 years, 11 months ago
B is the answer.
upvoted 2 times

aiwaai 4 years ago

Correct Answer: B

upvoted 1 times

paxjoshi 4 years ago

Yes, the correct answer is B.

upvoted 1 times

aiwaai 4 years ago

Correct Answer: B

upvoted 1 times

bigdo 4 years, 1 month ago

B D is for vulnerability scanning

upvoted 1 times

smart123 4 years, 1 month ago

The Answer is B

upvoted 1 times

Question #37 *Topic 1*

A company allows every employee to use Google Cloud Platform. Each department has a Google Group, with all department members as group members. If a department member creates a new project, all members of that department should automatically have read-only access to all new project resources. Members of any other department should not have access to the project. You need to configure this behavior.
What should you do to meet these requirements?

A. Create a Folder per department under the Organization. For each department's Folder, assign the Project Viewer role to the Google Group related to that department.

B. Create a Folder per department under the Organization. For each department's Folder, assign the Project Browser role to the Google Group related to that department.

C. Create a Project per department under the Organization. For each department's Project, assign the Project Viewer role to the Google Group related to that department.

D. Create a Project per department under the Organization. For each department's Project, assign the Project Browser role to the Google Group related to that department.

Correct Answer: *C*

ownez Highly Voted 👍 3 years, 11 months ago
Shouldn't it be A?

Project Browser has least permissions comparing to Project Viewer. The question is about have read-access to all new project resources.

roles/browser - Read access to browse the hierarchy for a project, including the folder, organization, and IAM policy. This role doesn't include permission to view resources in the project.

https://cloud.google.com/iam/docs/understanding-roles#project-roles
upvoted 21 times

singhjoga 3 years, 8 months ago
Correct, it is A. Project Browser does not have access to the resources inside the project, which is the requirement in the question.
upvoted 8 times

uiuiui Most Recent ⏲ 10 months ago
Selected Answer: A
A please
upvoted 1 times

IlDave 1 year, 6 months ago
Selected Answer: A
Create a Folder per department under the Organization. For each department's Folder, assign the Project Viewer role to the Google Group relatec to that department.
Grant viewer to the folder fits with automatically get permission on project creation
upvoted 2 times

mahi9 1 year, 6 months ago
Selected Answer: A
Create a Folder per department under the Organization. For each department's Folder, assign the Project Viewer role to the Google Group relatec to that department.
upvoted 1 times

Meyucho 1 year, 9 months ago
Selected Answer: A
Who voted C!?!??!?! The answer is A!!!!
upvoted 1 times

AwesomeGCP 1 year, 11 months ago
Selected Answer: A
Correct answer - A
https://cloud.google.com/resource-manager/docs/cloud-platform-resource-hierarchy
upvoted 1 times

piyush_1982 2 years, 1 month ago
Selected Answer: C

The correct answer is definitely C.
Let's divide the question into 2 parts:

1st: Role: Key requirement: all members of that department should automatically have read-only access to all new project resources.

> The project browser role only allows read access to browse the hierarchy for a project, including the folder, organization, and allow policy. This role doesn't include permission to view resources in the project.
Hence the options B and D are not relevant as they both are browser roles which DO NOT provide access to project resources.

2nd: Option A creates a Folder per department and C creates project per department.
However, Project viewer role is only applied at the project level.
Hence the correct answer is C which creates projects per department under organization .

upvoted 2 times

Meyucho 1 year, 9 months ago

But... if you dont have a folder per department.. where will be all new projects created by users???? you will have to manually edit permissions every time!!!! Using folders yu set the permitions once and then the only task you shoul do is to maintain the proper group assignment

upvoted 2 times

alvjtc 2 years, 1 month ago

Selected Answer: A

It's A, Project Viewer. Project Browser doesn't allow users to see resources, only find the project in the hierarchy.

upvoted 1 times

syllox 3 years, 4 months ago

It's A , browser is :
Read access to browse the hierarchy for a project, including the folder, organization, and IAM policy. This role doesn't include permission to view resources in the project.
https://cloud.google.com/iam/docs/understanding-roles#project-roles

upvoted 3 times

[Removed] 3 years, 4 months ago

either A or C because must be project viewer ,browser is not enough.https://cloud.google.com/iam/docs/understanding-roles

upvoted 1 times

[Removed] 3 years, 4 months ago

Why not A?

upvoted 1 times

desertlotus1211 3 years, 5 months ago

The answer is A:

https://stackoverflow.com/questions/54778596/whats-the-difference-between-project-browser-role-and-project-viewer-role-in-go#:~:text=8-,What's%20the%20difference%20between%20Project%20Browser%20role%20and,role%20in%20Google%20Cloud%20Platform&te=According%20to%20the%20console%20popup,read%20access%20to%20those%20resources.

upvoted 2 times

CloudTrip 3 years, 6 months ago

I think it's B. As the question says all members of that department should automatically have read-only access to all new project resources but browser will only provide the get, list permissions not read only permission so viewer seems to be more accurate here.

roles/browser
Read access to browse the hierarchy for a project, including the folder, organization, and IAM policy. This role doesn't include permission to view resources in the project.
resourcemanager.folders.get
resourcemanager.folders.list
resourcemanager.organizations.get
resourcemanager.projects.get
resourcemanager.projects.getIamPolicy
resourcemanager.projects.list

roles/viewer Viewer Permissions for read-only actions that do not affect state, such as viewing (but not modifying) existing resources or data.

upvoted 1 times

subhala 3 years, 9 months ago

Question says - If a department member creates a new project, all members of that department should automatically have read-only access to all new project resources. and @ownez provided documentation that says - browser role doesn't include perm to view resources in the project. Hen
B is the right answer.

upvoted 1 times

Fellipo 3 years, 9 months ago

A it´s OK

upvoted 2 times

[Removed] 3 years, 10 months ago

Ans - A

upvoted 2 times

cipher90 3 years, 11 months ago

Answer is B: "have read-only access to all new project resources." So it has to be in a folder to cascade the permissions to new projects carried.

upvoted 1 times

Meyucho 1 year, 9 months ago

If you do that the other members of the department can't access to the resourses.. just list the project in the folder

upvoted 1 times

Question #38 *Topic 1*

A customer's internal security team must manage its own encryption keys for encrypting data on Cloud Storage and decides to use customer-supplied encryption keys (CSEK).
How should the team complete this task?

A. Upload the encryption key to a Cloud Storage bucket, and then upload the object to the same bucket.

B. Use the gsutil command line tool to upload the object to Cloud Storage, and specify the location of the encryption key.

C. Generate an encryption key in the Google Cloud Platform Console, and upload an object to Cloud Storage using the specified key.

D. Encrypt the object, then use the gsutil command line tool or the Google Cloud Platform Console to upload the object to Cloud Storage.

Correct Answer: *D*
Reference:
https://cloud.google.com/storage/docs/encryption/customer-supplied-keys

DebasishLowes Highly Voted 3 years, 5 months ago
Ans : B. Because if you encrypt the object using CSEK, then you can't use google cloud console to upload the object.
upvoted 15 times

FatCharlie Highly Voted 3 years, 9 months ago
The fact is, both B & D would work. I lean towards B because it allows you to manage the file using GCP tools later as long as you keep that key around.

B is definitely incomplete though, as the boto file does need to be updated.
upvoted 7 times

gcpengineer 1 year, 3 months ago
it mentions u cant use console for CSEK
upvoted 1 times

3d9563b Most Recent 1 month, 1 week ago
Selected Answer: B
Using the gsutil command-line tool with the appropriate options to specify the CSEK during the upload process is the proper way to manage customer-supplied encryption keys for Cloud Storage. This ensures that the data is encrypted using the provided key without the key being store on Google's servers
upvoted 1 times

3d9563b 1 month, 1 week ago
Selected Answer: D
With Customer-Supplied Encryption Keys (CSEK), you handle the encryption of the data yourself and then upload the encrypted data to Cloud Storage, ensuring you provide the necessary encryption key when required for access control. This method ensures that you maintain control ove the encryption process and the security of your data.
upvoted 1 times

salamKvelas 3 months ago
`gcloud storage` you can point to a CSEK, but `gsutil` you can not
upvoted 1 times

shanwford 5 months ago
Selected Answer: B
Should be (B) - but IMHO "gsutil" is legacy tool, it works with "gcloud": gcloud storage cp SOURCE_DATA gs://BUCKET_NAME/OBJECT_NAME --encryption-key=YOUR_ENCRYPTION_KEY
upvoted 2 times

ppandher 10 months, 2 weeks ago
I have encrypt the object using 256 Encryption method, When I create a Bucket it gave me option of encryption as Google Managed Keys and Customer Managed keys but NO CSEK, I opted Google Managed as I do not have CMEK created, Now I create that Bucket.I upload my encrypted file to that bucket using Console, now the content of that file shows as Google managed not a CSEK.

To my understanding you need to generate the keys in console encrypt that object and then upload that way it will show on that object as encryption of CSEK.
Option B I opt now.
upvoted 1 times

mildi 1 year, 1 month ago
Answer D with removed or from console
D. Encrypt the object, then use the gsutil command line tool or the Google Cloud Platform Console to upload the object to Cloud Storage.

D. Encrypt the object, then use the gsutil command line tool

upvoted 1 times

twpower 1 year, 3 months ago

Selected Answer: B

Ans is B

upvoted 1 times

gcpengineer 1 year, 3 months ago

Selected Answer: B

B is the ans . https://cloud.google.com/storage/docs/encryption/customer-supplied-keys

upvoted 2 times

TQM__9MD 1 year, 4 months ago

Selected Answer: D

Object encryption is required. B does not encrypt objects.

upvoted 2 times

aashissh 1 year, 4 months ago

Selected Answer: D

To use customer-supplied encryption keys (CSEK) for encrypting data on Cloud Storage, the security team must encrypt the object first using the encryption key and then use the gsutil command line tool or the Google Cloud Platform Console to upload the object to Cloud Storage. Therefor the correct answer is:

D. Encrypt the object, then use the gsutil command line tool or the Google Cloud Platform Console to upload the object to Cloud Storage.

upvoted 2 times

gcpengineer 1 year, 3 months ago

it mentions u cant use console for CSEK

upvoted 1 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: B

https://cloud.google.com/storage/docs/encryption/customer-supplied-keys
Answer B

upvoted 2 times

GHOST1985 1 year, 11 months ago

Selected Answer: B

you can't use google cloud console to upload the object.
https://cloud.google.com/storage/docs/encryption/using-customer-supplied-keys#upload_with_your_encryption_key

upvoted 1 times

absipat 2 years, 2 months ago

D of course

upvoted 1 times

Aiffone 2 years, 2 months ago

I will go with D because encrypting the object before uploading means the cutomer manages thier own key.
A is not correct because its not a good practice to upload encryption key to storage object along with the encrypted object.
B is not correct because specifying the location of the encryption key does not change anything
C means Google manages the key.

upvoted 1 times

[Removed] 3 years, 4 months ago

CD are not right because Google Cloud Console does not support CSEK. must choose from A and B

upvoted 1 times

Question #39 *Topic 1*

A customer has 300 engineers. The company wants to grant different levels of access and efficiently manage IAM permissions between users in the development and production environment projects.
Which two steps should the company take to meet these requirements? (Choose two.)

A. Create a project with multiple VPC networks for each environment.

B. Create a folder for each development and production environment.

C. Create a Google Group for the Engineering team, and assign permissions at the folder level.

D. Create an Organizational Policy constraint for each folder environment.

E. Create projects for each environment, and grant IAM rights to each engineering user.

Correct Answer: *BD*

mozammil89 Highly Voted 4 years, 5 months ago

B and C should be correct...

upvoted 23 times

mahi9 Most Recent 1 year, 6 months ago

Selected Answer: BC

B and C are viable

upvoted 2 times

Meyucho 1 year, 9 months ago

Selected Answer: BC

Which Policy Constriaint allow to manage permission?!??!?! D is not an option. The answer is B and C

upvoted 2 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: BC

B and C are the correct answers!!

upvoted 2 times

danielklein09 2 years, 5 months ago

B is correct
But, if you make 1 group (by choosing option C) how you manage the permission for dev environment ? since you have only 1 group, you will off the same access for all 300 engineers (that are in that group) to dev and prod environment, so this will not answer the question: efficiently manag IAM permissions between users in the development and production environment projects

upvoted 4 times

Ksrp 2 years, 6 months ago

CE - A general recommendation is to have one project per application per environment. For example, if you have two applications, "app1" and "app2", each with a development and production environment, you would have four projects: app1-dev, app1-prod, app2-dev, app2-prod. This isolates the environments from each other, so changes to the development project do not accidentally impact production, and gives you better access control, since you can (for example) grant all developers access to development projects but restrict production access to your CI/CD pipeline. https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations

upvoted 1 times

Jane111 3 years, 4 months ago

A - no VPC required
B - yes - pre req
C - Yes
D - likely but C is first
E - not scalable/feasible/advisable

upvoted 2 times

DebasishLowes 3 years, 5 months ago

Ans : BC

upvoted 1 times

[Removed] 3 years, 10 months ago

Ans - BC

upvoted 1 times

CHECK666 3 years, 11 months ago

B,C is the answer.
Create a folder for each env and assign IAM policies to the group.

upvoted 2 times

MohitA 4 years ago

BC is the right answer, create folder for each env and assign IAM policies to group

upvoted 1 times

aiwaai 4 years ago

Correct Answer: CE

upvoted 1 times

> **aiwaai** 4 years ago
>
> made correction CE -> BC
>
> upvoted 2 times

xhova 4 years, 5 months ago

B&C

D does not help efficiently manage IAM. Effective IAM implies using groups.

upvoted 2 times

> **smart123** 4 years, 2 months ago
>
> Organization policy is used on resources and not the users. Hence option 'D' cannot be right.
>
> upvoted 2 times

jonclem 4 years, 5 months ago

I'd say B and D are correct

upvoted 1 times

Question #40 *Topic 1*

You want to evaluate your organization's Google Cloud instance for PCI compliance. You need to identify Google's inherent controls.
Which document should you review to find the information?

A. Google Cloud Platform: Customer Responsibility Matrix

B. PCI DSS Requirements and Security Assessment Procedures

C. PCI SSC Cloud Computing Guidelines

D. Product documentation for Compute Engine

Correct Answer: *C*
Reference:
https://cloud.google.com/solutions/pci-dss-compliance-in-gcp

3d9563b 1 month, 1 week ago
Selected Answer: A
The Customer Responsibility Matrix is the most relevant document for identifying Google's inherent controls related to PCI compliance, as it explicitly details the security controls managed by Google versus those managed by the customer.
upvoted 1 times

okhascorpio 6 months, 2 weeks ago
Selected Answer: A
Probably an outdated question, because there is a specific PCI DSS responsibility matrix available source:
https://cloud.google.com/security/compliance/pci-dss
but a close enough answer is A because it directly addresses Google's inherent controls while others don't.
upvoted 1 times

techdsmart 6 months, 3 weeks ago
but here controls isn't the same as responsibility? Don't understand how A is the answer since by controls we are referring this from a security an
compliance perspective i.e. security controls.
C is still the correct answer.
upvoted 1 times

rottzy 11 months, 2 weeks ago
answer is A, https://cloud.google.com/files/GCP_Client_Facing_Responsibility_Matrix_PCI_2018.pdf
upvoted 1 times

Xoxoo 11 months, 2 weeks ago
Selected Answer: A
To identify Google's inherent controls for PCI compliance, you should review:

A. Google Cloud Platform: Customer Responsibility Matrix

The Google Cloud Platform: Customer Responsibility Matrix provides information about the shared responsibility model between Google Cloud and the customer. It outlines which security controls are managed by Google and which are the customer's responsibility. This document will help you understand Google's inherent controls as they relate to PCI compliance.
upvoted 2 times

amanshin 1 year, 2 months ago
The correct answer is A. Google Cloud Platform: Customer Responsibility Matrix.

The Google Cloud Platform: Customer Responsibility Matrix (CRM) is a document that outlines the responsibilities of Google and its customers fo
PCI compliance. The CRM identifies the inherent controls that Google provides, which are the security controls that are built into Google Cloud Platform.

The PCI DSS Requirements and Security Assessment Procedures (SAQs) are a set of requirements that organizations must meet to be PCI compliant. The SAQs do not identify Google's inherent controls.

The PCI SSC Cloud Computing Guidelines are a set of guidelines that organizations can use to help them achieve PCI compliance when using clou
computing services. The guidelines do not identify Google's inherent controls.

The product documentation for Compute Engine is a document that provides information about the features and capabilities of Compute Engine
The documentation does not identify Google's inherent controls.
upvoted 1 times

gcpengineer 1 year, 3 months ago
Selected Answer: C
C is the ans

upvoted 2 times

gcpengineer 1 year, 3 months ago

Selected Answer: B

B is the ans. as the pci-dss req in gcp

upvoted 1 times

gcpengineer 1 year, 3 months ago

C is the ans

upvoted 1 times

aashissh 1 year, 4 months ago

Selected Answer: A

The answer is A. Google Cloud Platform: Customer Responsibility Matrix. This document outlines the responsibilities of both the customer and Google for securing the cloud environment and is an important resource for understanding Google's inherent controls for PCI compliance. The PC DSS Requirements and Security Assessment Procedures and the PCI SSC Cloud Computing Guidelines are both helpful resources for understandi the PCI compliance requirements, but they do not provide information on Google's specific inherent controls. The product documentation for Compute Engine is focused on the technical aspects of using that service and is unlikely to provide a comprehensive overview of Google's inhere controls.

upvoted 3 times

1explorer 1 year, 5 months ago

https://cloud.google.com/architecture/pci-dss-compliance-in-gcp
B is correct answer

upvoted 3 times

tailesley 1 year, 6 months ago

It is B:: The PCI DSS Requirements and Security Assessment Procedures is the document that outlines the specific requirements for PCI complianc It is created and maintained by the Payment Card Industry Security Standards Council (PCI SSC), which is the organization responsible for establishing and enforcing security standards for the payment card industry. This document is used by auditors to evaluate the security of an organization's payment card systems and processes.

While the other options may provide information about Google's security controls and the customer's responsibilities for security, they do not provide the specific requirements for PCI compliance that the PCI DSS document does.

upvoted 3 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: A

A. Google Cloud Platform: Customer Responsibility Matrix

upvoted 1 times

tangac 1 year, 12 months ago

Selected Answer: A

https://services.google.com/fh/files/misc/gcp_pci_shared_responsibility_matrix_aug_2021.pdf

upvoted 2 times

Question #41 *Topic 1*

Your company runs a website that will store PII on Google Cloud Platform. To comply with data privacy regulations, this data can only be stored for a specific amount of time and must be fully deleted after this specific period. Data that has not yet reached the time period should not be deleted. You want to automate the process of complying with this regulation.
What should you do?

A. Store the data in a single Persistent Disk, and delete the disk at expiration time.

B. Store the data in a single BigQuery table and set the appropriate table expiration time.

C. Store the data in a single Cloud Storage bucket and configure the bucket's Time to Live.

D. Store the data in a single BigTable table and set an expiration time on the column families.

Correct Answer: *B*

KILLMAD Highly Voted 4 years, 5 months ago

I believe the Answer is C not B.

This isn't data which needs to be analyzed, so I don't understand why would it be stored in BQ when having data stored in GCS seems much mor
reasonable.

I think the only thing about answer C which throws me off is the fact that they don't mention object life cycle management

upvoted 14 times

mozammil89 4 years, 5 months ago

Answer C is correct. The TTL is common use case of Cloud Storage life cycle management. Here is what GCP says:

"To support common use cases like setting a Time to Live (TTL) for objects, retaining noncurrent versions of objects, or "downgrading" storage
classes of objects to help manage costs, Cloud Storage offers the Object Lifecycle Management feature. This page describes the feature as we
as the options available when using it. To learn how to enable Object Lifecycle Management, and for examples of lifecycle policies, see Managing Lifecycles."

https://cloud.google.com/storage/docs/lifecycle

upvoted 7 times

PleeO 3 months, 2 weeks ago

This answer is still valid till 2024

upvoted 1 times

trashbox Most Recent 4 months ago

Selected Answer: C

Bucket lock and TTL are the key features of Cloud Storage.

upvoted 1 times

Bypoo 6 months, 2 weeks ago

Selected Answer: C

Cloud Storage life cycle management

upvoted 1 times

Echizen06 1 year ago

Selected Answer: C

Answer is C

upvoted 2 times

cyberpunk21 1 year ago

B is correct, all forgot this "Data that has not yet reached the time period should not be deleted." from question this means data is keep on updating if we enforce TTL for a bucker the whole bucket will be deleted including updated data, so with Big query we do updating using pipelin
jobs and delete data using expiration time

upvoted 1 times

mahi9 1 year, 6 months ago

Selected Answer: C

store it in a bucket for TTL

upvoted 2 times

PST21 1 year, 8 months ago

CS does not delete promptly , hence BQ as it is sensitive data

upvoted 1 times

csrazdan 1 year, 9 months ago

Selected Answer: B

Life Cycle Management for Cloud storage is used to manage the Storage class to save cost. For data management, you have set retention time or the bucket. I will opt for B as the correct answer.

upvoted 1 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: C

Correct Answer: C

upvoted 2 times

giovy_82 2 years ago

I would go for C, but all the 4 answers are in my opinion incomplete. all of them say "single" bucket or table, which means that if different dated rows/elements are stored in the same bucket or table, they will expire together and be deleted probably before their real expiration time. so i expected to see partitioning or multiple bucket.

upvoted 2 times

mynk29 2 years, 6 months ago

Outdated question again- should be bucket locks now.

upvoted 1 times

DebasishLowes 3 years, 5 months ago

Ans : C

upvoted 2 times

[Removed] 3 years, 10 months ago

Ans - C

upvoted 4 times

aiwaai 4 years ago

Correct Answer: C

upvoted 3 times

Ganshank 4 years, 3 months ago

The answers need to be worded better.
If we're taking the terms literally as specified in the options, then C cannot be the correction answer since there's no Time to Live configuration fo a GCS bucket, only Lifecycle Policy.
With BigQuery, there is no row-level expiration, although we could create this behavior using Partitioned Tables. So this could be a potential answer.
D - it is possible to simulate cell-level TTL (https://cloud.google.com/bigtable/docs/gc-cell-level), so this too could be a potential answer, especia when different cells need different TTLs.
Betweem B & D, BigQuery follows a pay-as-you-go model and its storage costs are comparable to GCS storage costs. So this would be the more appropriate solution.

upvoted 3 times

> **smart123** 4 years, 2 months ago
>
> The Buckets do have "Time to Live" feature.
> https://cloud.google.com/storage/docs/lifecycle
>
> Hence 'C' is the answer
>
> upvoted 4 times

jonclem 4 years, 5 months ago

I believe B is correct.

Setting a TTL of 14 days on the bucket via LifeCycle will not cause the bucket itself to be deleted after 14 days, instead it will cause each object uploaded to that bucket to be deleted 14 days after it was created

upvoted 3 times

> **xhova** 4 years, 5 months ago
>
> Answer is C. You dont need the bucket to be deleted, you need the PII data stored to be deleted.
>
> upvoted 6 times

Question #42 *Topic 1*

A DevOps team will create a new container to run on Google Kubernetes Engine. As the application will be internet-facing, they want to minimize the attack surface of the container.
What should they do?

A. Use Cloud Build to build the container images.

B. Build small containers using small base images.

C. Delete non-used versions from Container Registry.

D. Use a Continuous Delivery tool to deploy the application.

Correct Answer: *D*
Reference:
https://cloud.google.com/solutions/best-practices-for-building-containers

xhova Highly Voted 4 years, 5 months ago

Ans is B

Small containers usually have a smaller attack surface as compared to containers that use large base images.

https://cloud.google.com/blog/products/gcp/kubernetes-best-practices-how-and-why-to-build-small-container-images

upvoted 31 times

smart123 4 years, 1 month ago

I agree

upvoted 2 times

3d9563b Most Recent 1 month, 1 week ago

Selected Answer: B

Building small containers using minimal and well-maintained base images directly reduces the attack surface and improves the security posture o your containers when they are deployed on GKE.

upvoted 1 times

okhascorpio 6 months, 2 weeks ago

Selected Answer: B

the correct answer is having as few tools in your image as possible, Source: Remove unnecessary tools https://cloud.google.com/architecture/bes practices-for-building-containers?hl=en
I guess it can be achieved by option "B" building a small container from a small source image.

upvoted 1 times

Afe3saa7 6 months, 3 weeks ago

Selected Answer: B

A. Use Cloud Build to build the container images.
Will give you the tools to build an image but not ensure any risk reduction

B. Build small containers using small base images.
Images with a smaller footprint, stripped of all binaries/libraries/functions that are not used will make it harder for an attacker to find leverage to move laterally or vertically, hence >>reducing the attack/risk surface<< for the image.

C. Delete non-used versions from Container Registry.
Non-used images are not running live and hence are not exploitable. Removing non-used images from the registry will not reduce the attack surface of the running application.

D. Use a Continuous Delivery tool to deploy the application.
Same as A.

upvoted 1 times

Xoxoo 11 months, 2 weeks ago

Selected Answer: B

To minimize the attack surface of a container that will run on Google Kubernetes Engine and be internet-facing, the DevOps team should:

B. Build small containers using small base images.

Building small containers using minimal base images reduces the attack surface by eliminating unnecessary software and dependencies, which ca potentially contain vulnerabilities. This approach enhances security and reduces the risk of potential attacks. Using small base images, such as Alpine Linux or distroless images, is a best practice for container security.

upvoted 3 times

civilizador 1 year, 1 month ago

Answer is B, because this GCP exam, the GCP docs are always source of truth even though you might not be agree with them occasionally but eve
if you are not agree you need to choose the answer proposed in GCP docs as the best practice.
Here is the link to google official best practices for building containers. and here is the snippet regarding this particular question:
https://cloud.google.com/architecture/best-practices-for-building-containers#build-the-smallest-image-possible

Build the smallest image possible
Building a smaller image offers advantages such as faster upload and download times, which is especially important for the cold start time of a po
in Kubernetes: the smaller the image, the faster the node can download it. However, building a small image can be difficult because you might inadvertently include build dependencies or unoptimized layers in your final image.

upvoted 2 times

[Removed] 1 year, 1 month ago

Selected Answer: B

"B"
For smaller attacker surface, use smaller images by removing any unnecessary tools/software from the image.

https://cloud.google.com/solutions/best-practices-for-building-containers

upvoted 2 times

alleinallein 1 year, 5 months ago

Selected Answer: C

Importance: MEDIUM

To protect your apps from attackers, try to reduce the attack surface of your app by removing any unnecessary tools.

https://cloud.google.com/architecture/best-practices-for-building-containers

upvoted 2 times

adb4007 9 months, 1 week ago

So build a small image is the answer, not ?

upvoted 1 times

mahi9 1 year, 6 months ago

Selected Answer: C

it is viable

upvoted 1 times

rotorclear 1 year, 10 months ago

Selected Answer: B

B definitely

upvoted 1 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

zellck 1 year, 11 months ago

Selected Answer: B

B is the answer.

upvoted 1 times

jitu028 1 year, 11 months ago

Ans is B - https://cloud.google.com/blog/products/gcp/kubernetes-best-practices-how-and-why-to-build-small-container-images

Security and vulnerabilities
Aside from performance, there are significant security benefits from using smaller containers. Small containers usually have a smaller attack surfac
as compared to containers that use large base images.

upvoted 3 times

giovy_82 2 years ago

Selected Answer: B

the only answer that will really reduce attack surface while exposing apps to internet is B, small containers (e.g. single web page?)

upvoted 3 times

Medofree 2 years, 4 months ago

B. Because you will have less programs in the image thus less vulnerabilities

upvoted 1 times

lxs 2 years, 9 months ago

Selected Answer: C

A. Use Cloud Build to build the container images.
If you build a container using Cloud Build or not the surface is the same

B. Build small containers using small base images.
It is indeed best practice, but I doubt if small base images can reduce the surface. It is still the same app version with the same vulnerabilities etc.
C. Delete non-used versions from Container Registry.
Unused, historical versions are additional attack surface. attacker can exploit old, unpatched image which indeed the surface extention.
D. Use a Continuous Delivery tool to deploy the application.
This is just a method of image delivery. The app is the same.

upvoted 3 times

Afe3saa7 6 months, 3 weeks ago

non-used images in containter registry are as they suggest not running live, hence are not exploitable. deleting images in the registry will not change the attack surface of the mentioned image.

upvoted 1 times

DebasishLowes 3 years, 5 months ago

Ans : B. Small the base image there is less vulnerability and less chance of attack.

upvoted 2 times

Question #43 *Topic 1*

While migrating your organization's infrastructure to GCP, a large number of users will need to access GCP Console. The Identity Management team already has a well-established way to manage your users and want to keep using your existing Active Directory or LDAP server along with the existing SSO password.
What should you do?

A. Manually synchronize the data in Google domain with your existing Active Directory or LDAP server.

B. Use Google Cloud Directory Sync to synchronize the data in Google domain with your existing Active Directory or LDAP server.

C. Users sign in directly to the GCP Console using the credentials from your on-premises Kerberos compliant identity provider.

D. Users sign in using OpenID (OIDC) compatible IdP, receive an authentication token, then use that token to log in to the GCP Console.

Correct Answer: *B*
Reference:
https://cloud.google.com/blog/products/identity-security/using-your-existing-identity-management-system-with-google-cloud-platform

sudarchary Highly Voted 2 years, 7 months ago
Selected Answer: B
https://cloud.google.com/architecture/identity/federating-gcp-with-active-directory-configuring-single-sign-on
upvoted 7 times

DebasishLowes Highly Voted 3 years, 5 months ago
Ans : B
upvoted 5 times

dbf0a72 Most Recent 8 months ago
Selected Answer: B
https://cloud.google.com/architecture/identity/federating-gcp-with-active-directory-configuring-single-sign-on
upvoted 1 times

AwesomeGCP 1 year, 11 months ago
Selected Answer: B
https://cloud.google.com/architecture/identity/federating-gcp-with-active-directory-configuring-single-sign-on
upvoted 2 times

absipat 2 years, 2 months ago
B of course
upvoted 2 times

ThisisJohn 2 years, 8 months ago
Selected Answer: D
My vote goes for D.

From the blog post linked below " users' passwords are not synchronized by default. Only the identities are synchronized, unless you make an explicit choice to synchronize passwords (which is not a best practice and should be avoided)".

Also, from GCP documentation "Authenticating with OIDC and AD FS" https://cloud.google.com/anthos/clusters/docs/on-prem/1.6/how-to/oidc-adfs

Blog post quoted above https://cloud.google.com/blog/products/identity-security/using-your-existing-identity-management-system-with-googl cloud-platform
upvoted 1 times

rr4444 2 years, 8 months ago
D sounds nice, but the user doesn't "use" the token.... that's used in the integration with Cloud Identity. So answer must be B, GCDS
upvoted 3 times

[Removed] 3 years, 10 months ago
Ans - B
upvoted 4 times

saurabh1805 3 years, 10 months ago
B is correct answer here.
upvoted 4 times

Question #44 *Topic 1*

Your company is using GSuite and has developed an application meant for internal usage on Google App Engine. You need to make sure that an external user cannot gain access to the application even when an employee's password has been compromised.
What should you do?

A. Enforce 2-factor authentication in GSuite for all users.

B. Configure Cloud Identity-Aware Proxy for the App Engine Application.

C. Provision user passwords using GSuite Password Sync.

D. Configure Cloud VPN between your private network and GCP.

Correct Answer: *D*

rafaelc Highly Voted 👍 4 years, 5 months ago

A. Enforce 2-factor authentication in GSuite for all users.

upvoted 21 times

johnsm Highly Voted 👍 3 years, 1 month ago

Correct answer is A. Well explained here: https://docs.google.com/document/d/11o3e14tyhnT7w45Q8-r9ZmTAfj2WUNUpJPZImrxm_F4/edit?usp=sharing found some other answers for other questions in this site as well.

upvoted 8 times

Oujay Most Recent ⊙ 2 months ago

Selected Answer: B

A Cloud VPN creates a secure tunnel between your network and GCP, but it wouldn't restrict access based on individual user identities.

upvoted 1 times

Oujay 2 months ago

2FA adds an extra layer of security, but if an external user has both the password and the second factor (e.g., a verification code), they might still gain access.
So my answer is B. All external users will be blocked with the right authentication or not

upvoted 1 times

dbf0a72 8 months ago

Selected Answer: A

A is the answer.

upvoted 1 times

raj117 1 year, 1 month ago

Right Answer is A

upvoted 2 times

SMB2022 1 year, 1 month ago

Correct Answer A

upvoted 2 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: A

A is the answer.

upvoted 3 times

sudarchary 2 years, 7 months ago

Selected Answer: A

https://support.google.com/a/answer/175197?hl=en

upvoted 2 times

Jane111 3 years, 4 months ago

Shouldn't it be
B. Configure Cloud Identity-Aware Proxy for the App Engine Application.
identity based app access

upvoted 4 times

[Removed] 1 year, 1 month ago

I was thinking the same thing. Turns out IAP ensures security by enforcing 2FA. So at the end of the day, 2FA is the real solution.
2FA without IAP would still address the risk. IAP without 2FA might not.
https://cloud.google.com/iap/docs/configuring-reauth#supported_reauthentication_methods

upvoted 2 times

desertlotus1211 3 years, 5 months ago

The key is external user. Best practice is to have internal users/datacenter connect via VPN for security purpose, correct? External users will try to connect via Internet - they still cannot reach the app engine even if they have a users' password because a VPN connection is need to reach the resource. MA will work IF the external user has VPN access... But I think D is what they're looking for based on the question....

upvoted 3 times

mynk29 2 years, 6 months ago

Agree but there is no mention that external user doesnt have internal network access too. A is better option as it covers both scenarios.

upvoted 2 times

DebasishLowes 3 years, 5 months ago

Ans : A. When passwords is compromised, enforcing 2 factor authentication is the best way to prevent non authorized users.

upvoted 2 times

soukumar369 3 years, 8 months ago

Enforcing 2-factor authentication can save an employee's password has been compromised

upvoted 2 times

soukumar369 3 years, 8 months ago

Enforce 2-factor authentication safe employee, when an employee's password has been compromised.

upvoted 2 times

[Removed] 3 years, 10 months ago

Ans - A

upvoted 2 times

subhala 3 years, 10 months ago

If you limit your GCP VPC to only private access (no resources having external IP), and have VPN. then inspite of having any creds, external folks cannot access the resources.

upvoted 2 times

Cloudy_Apple_Juice 3 years, 10 months ago

They can if they login from inside Org - So A is the only correct asnwer

upvoted 1 times

soukumar369 3 years, 9 months ago

I'm also thinking the same.

upvoted 2 times

passtest100 3 years, 11 months ago

should be B

upvoted 2 times

Question #45

Topic 1

A large financial institution is moving its Big Data analytics to Google Cloud Platform. They want to have maximum control over the encryption process of data stored at rest in BigQuery.
What technique should the institution use?

A. Use Cloud Storage as a federated Data Source.

B. Use a Cloud Hardware Security Module (Cloud HSM).

C. Customer-managed encryption keys (CMEK).

D. Customer-supplied encryption keys (CSEK).

Correct Answer: *C*
Reference:
https://cloud.google.com/bigquery/docs/encryption-at-rest

Ganshank Highly Voted 4 years, 3 months ago
CSEK is only supported in Google Cloud Storage and Compute Engine, therefore D cannot be the right answer.
Ideally, it would be client-side encryption, with BigQuery providing another round of encryption of the encrypted data - https://cloud.google.com/bigquery/docs/encryption-at-rest#client_side_encryption, but since that is not one of the options, we can go with C as the next best option.
upvoted 19 times

smart123 4 years, 2 months ago
Option 'C' is correct. Option 'D' is not correct as CSEK a feature in Google Cloud Storage and Google Compute Engine only.
upvoted 5 times

crazycosmos Most Recent 1 month ago
Selected Answer: D
I prefer D for max control.
upvoted 1 times

SQLbox 1 month, 1 week ago
Correct answer is D

D. Customer-supplied encryption keys (CSEK).

Here's an explanation of why CSEK is the best choice and a brief review of the other options:

Customer-supplied encryption keys (CSEK): CSEK allows the institution to manage their own encryption keys and supply these keys to Google Cloud Platform when needed. This provides maximum control over the encryption process because the institution retains possession of the encryption keys and can rotate, revoke, or replace them as desired.
upvoted 1 times

Ishu_awsguy 1 year, 3 months ago
Why not Cloud HSM ?
Maximum control over keys
upvoted 1 times

Ishu_awsguy 1 year, 3 months ago
Sorry
From HSM the keys become customer supplied encryption keys which are not supported.
Ans is Customer managed encryptipn keys
upvoted 1 times

AwesomeGCP 1 year, 11 months ago
Selected Answer: C
C. Customer-managed encryption keys (CMEK).
upvoted 2 times

DebasishLowes 3 years, 5 months ago
Ans : C
upvoted 2 times

Aniyadu 3 years, 8 months ago
I feel C is the right answer. if customer wants to manage the keys from on-premises then D would be correct.
upvoted 3 times

[Removed] 3 years, 10 months ago

Ans - C

upvoted 3 times

saurabh1805 3 years, 10 months ago

C is correct answer as CSEK is not available for big query.

upvoted 2 times

MohitA 4 years ago

C is the right answer as CSEC is only available for CS and CE's

upvoted 1 times

aiwaai 4 years ago

Correct Answer: C

upvoted 2 times

ArizonaClassics 4 years, 1 month ago

C is the RIGHT ONE!!!

If you want to manage the key encryption keys used for your data at rest, instead of having Google manage the keys, use Cloud Key Managemen Service to manage your keys. This scenario is known as customer-managed encryption keys (CMEK).
https://cloud.google.com/bigquery/docs/encryption-at-rest

upvoted 2 times

> **ArizonaClassics** 4 years ago
>
> ALSO READ : https://cloud.google.com/bigquery/docs/customer-managed-encryption
>
> upvoted 2 times

ranjeetpatil 4 years, 2 months ago

Ans is C. BigQuery does not support CSEK. https://cloud.google.com/security/encryption-at-rest. https://cloud.google.com/security/encryption-at rest

upvoted 4 times

srinidutt 4 years, 3 months ago

I also feeel D is right

upvoted 1 times

xhova 4 years, 5 months ago

Answer is D. For max control you don't want to store the Key with Google.

upvoted 3 times

jonclem 4 years, 5 months ago

For maximum control surely D is the correct answer.
CSEK:
https://cloud.google.com/security/encryption-at-rest/customer-supplied-encryption-keys

CMEK
https://cloud.google.com/bigquery/docs/encryption-at-rest

upvoted 2 times

Question #46 *Topic 1*

A company is deploying their application on Google Cloud Platform. Company policy requires long-term data to be stored using a solution that can automatically replicate data over at least two geographic places.
Which Storage solution are they allowed to use?

A. Cloud Bigtable

B. Cloud BigQuery

C. Compute Engine SSD Disk

D. Compute Engine Persistent Disk

Correct Answer: *B*
Reference:
https://cloud.google.com/bigquery/docs/locations

ronron89 Highly Voted 3 years, 8 months ago
https://cloud.google.com/bigquery#:~:text=BigQuery%20transparently%20and%20automatically%20provides,charge%20and%20no%20additiona
%20setup.&text=BigQuery%20also%20provides%20ODBC%20and,interact%20with%20its%20powerful%20engine.

Answer is B.

BigQuery transparently and automatically provides highly durable, replicated storage in multiple locations and high availability with no extra charge and no additional setup.

@xhova: https://cloud.google.com/bigquery-transfer/docs/locations
What it mentions here is once you create a replication. YOu cannot change a location. Here the question is about high availability. synchronous replication.
upvoted 15 times

Arad 2 years, 9 months ago
Correct answer is A.
B is not correct because: "BigQuery does not automatically provide a backup or replica of your data in another geographic region."
https://cloud.google.com/bigquery/docs/availability
upvoted 6 times

mynk29 2 years, 6 months ago
"In either case, BigQuery automatically stores copies of your data in two different Google Cloud zones within the selected location."

your link
upvoted 4 times

mistryminded 2 years, 9 months ago
Correct answer is B.

BQ: https://cloud.google.com/bigquery-transfer/docs/locations#multi-regional-locations and https://cloud.google.com/bigquery-transfer/docs/locations#colocation_required

Bigtable: https://cloud.google.com/bigtable/docs/locations

PS: To people that are only commenting an answer, please provide a valid source to back your answers. This is a community driven forum and just spamming with wrong answers affects all of us.
upvoted 7 times

ryumoe Most Recent 2 months, 1 week ago
Answer is D, becasue:

A. Cloud Bigtable: This is a NoSQL database service, not designed for long-term data storage with automatic geographic replication.
B. Cloud BigQuery: This is a data warehouse service, excellent for analyzing data, but it doesn't inherently replicate data for disaster recovery.
C. Compute Engine SSD Disk: These are local disks attached to virtual machines, not designed for long-term storage or automatic replication.
upvoted 1 times

nccdebug 6 months, 2 weeks ago
BigQuery automatically stores copies of your data in two different Google Cloud zones within a single region in the selected location.
https://cloud.google.com/bigquery/docs/locations
upvoted 1 times

adb4007 9 months ago
In my opinion the key word is "automatic" because BigQuery and BigeTable are by default store on one zone for a piece of data (no replication)
Withe BigTable replication is automatic : https://cloud.google.com/bigtable/docs/replication-overview and copy dataset on Bigquery is not

automatic https://cloud.google.com/bigquery/docs/managing-datasets#copy-datasets I go to A

upvoted 1 times

uiuiui 10 months ago

Selected Answer: D

this is geographic, not region, then the correct ans is D

upvoted 1 times

civilizador 1 year, 1 month ago

Answer is A - Cloud Bigtable.
Cloud Bigtable - Replication: This page provides a detailed overview of how Cloud Bigtable uses replication to increase the availability and durability of your data.

Cloud BigQuery: From the BigQuery product description, you can see that it is mainly focused on analyzing data and does not mention geograph replication of data as a feature.

Compute Engine Disks: The documentation for Compute Engine Disks explains that they are zonal resources, meaning they are replicated within a single zone, but not across multiple zones or regions.

upvoted 1 times

megalucio 1 year, 1 month ago

Selected Answer: A

Correct one is A, as BigQuery does not provide replication but multi location storage which is different

upvoted 1 times

Ishu_awsguy 1 year, 2 months ago

I am drifting towards D
Regional persistent disk are safe from zonal failures.
The question mentions different geo places (not regions) .
So if zone seperation is done in 1 google region and we use regional persistent disk , the data will be safe from failure.
Also why would someone move their DR to BQ ? persistent disk make more sense to me

upvoted 1 times

Ishu_awsguy 1 year, 3 months ago

Point not to be confused ,
Even with BQ multi region , data s stores in different ones in 1 region not different geographic regions.

The question asks " different geographic places " which means essentially seperate zone storage will work.
hence answer is B (Big query) either single region or multi region .
Both suffice

upvoted 1 times

Ishu_awsguy 1 year, 3 months ago

--- Typo correction ---
Point not to be confused ,
Even with BQ multi region , data is stored in different zones in 1 region & not different geographic regions.

The question asks " different geographic places " which means essentially separate zone storage will work.
hence answer is B (Big query) either single region or multi region .
Both suffice

upvoted 1 times

deony 1 year, 3 months ago

I think answer is B
First of reason is long-term data solution, it's suitable for Cloud Storage and BigQuery
Second is that BigQuery dataset is placed to multi-region that means that two or more regions.

upvoted 1 times

Ric350 1 year, 5 months ago

The answer is definitely A. Here's why: https://cloud.google.com/bigtable/docs/replication-overview#how-it-works
Replication for Cloud Bigtable lets you increase the availability and durability of your data by copying it across multiple regions or multiple zones within the same region. You can also isolate workloads by routing different types of requests to different clusters.

BQ does not do cross-region replication. The blue highlighted note in the two links below clearly says the following: "Selecting a multi-region location does NOT provide cross-region replication NOR regional redundancy. Data will be stored in a single region within the geographic location."
https://cloud.google.com/bigquery/docs/reliability-disaster#availability_and_durability
https://cloud.google.com/bigquery/docs/locations#multi-regions

upvoted 4 times

sameer2803 1 year, 6 months ago

Answer is A.
the below statement is from the google cloud documentation. https://cloud.google.com/bigquery/docs/reliability-disaster
BigQuery does not automatically provide a backup or replica of your data in another geographic region

upvoted 3 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: B

B. Cloud BigQuery

upvoted 1 times

giovy_82 2 years ago

Selected Answer: B

I was about to select D, BUT:
- the question says "long term data" -> which makes me think about BQ
- the replication of persistent disk is between different ZONES, but the question says "different geo location" -> which means different regions (if you look at the zone distribution, different zones in same region are located in the same datacenter)

but I still have doubt since the application data are not supposed to be stored in BQ , unless it is for analytics and so on. GCS would have been th best choice, but in absence of this, probably B is the 1st choice.

upvoted 3 times

Table2022 1 year, 10 months ago

Thank God we have you giovy_82, very good explanation.

upvoted 2 times

piyush_1982 2 years, 1 month ago

Selected Answer: A

https://cloud.google.com/bigquery/docs/availability#availability_and_durability

As per the link above BigQuery does not automatically provide a backup or replica of your data in another geographic region. It only stores copie of data in two different Google Cloud zones within the selected location.

Reading through the link https://cloud.google.com/bigtable/docs/replication-overview
It states that the Bigtable replicates any changes to your data automatically within a region or multi-region.

upvoted 2 times

Medofree 2 years, 4 months ago

Selected Answer: D

The answer is D.

Do not forget the context, we are talking here about an Application and its storage, we are not talking about high throughput low latency databa (Bigtable) or Analytics Database (BigQuery). On the other way SSDs are physically attached to the VM.

upvoted 2 times

sudarchary 2 years, 7 months ago

Selected Answer: B

Multiregion

upvoted 1 times

Question #47 *Topic 1*

A large e-retailer is moving to Google Cloud Platform with its ecommerce website. The company wants to ensure payment information is encrypted between the customer's browser and GCP when the customers checkout online.
What should they do?

A. Configure an SSL Certificate on an L7 Load Balancer and require encryption.

B. Configure an SSL Certificate on a Network TCP Load Balancer and require encryption.

C. Configure the firewall to allow inbound traffic on port 443, and block all other inbound traffic.

D. Configure the firewall to allow outbound traffic on port 443, and block all other outbound traffic.

Correct Answer: *A*

ESP_SAP Highly Voted 👍 3 years, 9 months ago
Correct Answer is (A):

he type of traffic that you need your load balancer to handle is another factor in determining which load balancer to use:

For HTTP and HTTPS traffic, use:
External HTTP(S) Load Balancing

https://cloud.google.com/load-balancing/docs/load-balancing-overview#external_versus_internal_load_balancing
upvoted 11 times

fandyadam Most Recent ⊙ 9 months, 3 weeks ago
Selected Answer: A
upvoted 2 times

pedrojorge 1 year, 7 months ago
Selected Answer: A
A is right
upvoted 2 times

[Removed] 3 years, 10 months ago
Ans - A
upvoted 2 times

CHECK666 3 years, 11 months ago
A is the answer, SSL certificate on L7 layer LoadBlanacer
upvoted 3 times

ArizonaClassics 4 years, 1 month ago
A is the correct one. the question is to see if you understand difference between Layer 7 vs Layer 4 protocols.
upvoted 2 times

smart123 4 years, 1 month ago
Option 'A' is the correct answer.
upvoted 1 times

srinidutt 4 years, 3 months ago
A is right
upvoted 1 times

Question #48 *Topic 1*

Applications often require access to `secrets` - small pieces of sensitive data at build or run time. The administrator managing these secrets on GCP wants to keep a track of `who did what, where, and when?` within their GCP projects.
Which two log streams would provide the information that the administrator is looking for? (Choose two.)

A. Admin Activity logs

B. System Event logs

C. Data Access logs

D. VPC Flow logs

E. Agent logs

Correct Answer: *AC*
Reference:
https://cloud.google.com/kms/docs/secret-management

Ganshank Highly Voted 👍 4 years, 3 months ago
Agreed AC.
https://cloud.google.com/secret-manager/docs/audit-logging
upvoted 13 times

ArizonaClassics Most Recent 🕘 1 year ago
AC: Read https://cloud.google.com/logging/docs/audit#admin-activity
upvoted 2 times

[Removed] 1 year, 1 month ago
Selected Answer: AC
A, C.
https://cloud.google.com/secret-manager/docs/audit-logging#available-logs
upvoted 3 times

AwesomeGCP 1 year, 11 months ago
Selected Answer: AC
A. Admin Activity logs
C. Data Access logs
upvoted 2 times

DebasishLowes 3 years, 6 months ago
Ans AC
upvoted 4 times

[Removed] 3 years, 10 months ago
Ans - AC
upvoted 2 times

CHECK666 3 years, 11 months ago
AC is the answer.
Admin Access Logs and Data Access Logs
upvoted 3 times

smart123 4 years, 1 month ago
Yes 'A & C' are the right answers.
upvoted 2 times

Question #49 *Topic 1*

You are in charge of migrating a legacy application from your company datacenters to GCP before the current maintenance contract expires. You do not know what ports the application is using and no documentation is available for you to check. You want to complete the migration without putting your environment at risk.
What should you do?

A. Migrate the application into an isolated project using a ג€Lift & Shiftג€ approach. Enable all internal TCP traffic using VPC Firewall rules. Use VPC Flow logs to determine what traffic should be allowed for the application to work properly.

B. Migrate the application into an isolated project using a ג€Lift & Shiftג€ approach in a custom network. Disable all traffic within the VPC and look at the Firewall logs to determine what traffic should be allowed for the application to work properly.

C. Refactor the application into a micro-services architecture in a GKE cluster. Disable all traffic from outside the cluster using Firewall Rules. Use VPC Flow logs to determine what traffic should be allowed for the application to work properly.

D. Refactor the application into a micro-services architecture hosted in Cloud Functions in an isolated project. Disable all traffic from outside your project using Firewall Rules. Use VPC Flow logs to determine what traffic should be allowed for the application to work properly.

Correct Answer: *C*

rafaelc Highly Voted 4 years, 5 months ago

A or B. Leaning towards A
You have a deadline you cannot develop a new app so you have to lift and shift.

upvoted 20 times

mynk29 2 years, 6 months ago

Agree "Disable all traffic within the VPC and look at the Firewall logs to determine what traffic should be allowed for the application to work properly." if you disable all the VPC traffic there will be nothing to look into firewall logs.

upvoted 7 times

xhova 4 years, 5 months ago

Answer is A.. You need VPC Flow Logs not "Firewall logs" stated in B

upvoted 13 times

smart123 4 years, 1 month ago

I agree.

upvoted 2 times

Table2022 1 year, 10 months ago

xhova, you got it right!

upvoted 3 times

cskhachane Most Recent 6 months, 1 week ago

Option C:

upvoted 1 times

okhascorpio 6 months, 2 weeks ago

Selected Answer: A

B is not correct because Disabling all traffic within the VPC is too restrictive and hinders even initial testing. Analyzing firewall logs without any initial connectivity wouldn't be feasible.

upvoted 2 times

Xoxoo 11 months, 2 weeks ago

Selected Answer: A

Option B, C, and D involve making significant architectural changes (refactoring into microservices or using Cloud Functions) and disabling traffic, which might introduce complexities and risks. These options are more suitable when you have a better understanding of the application's requirements and can make informed decisions about its architecture and network policies. In your current scenario, option A provides a safe starting point for the migration process while you gather more information about the application's behavior.

upvoted 3 times

ArizonaClassics 11 months, 3 weeks ago

B. This option is similar to the first one but is more secure initially. The application is also migrated using a "Lift & Shift" approach. However, instead of enabling all internal TCP traffic, all traffic within the VPC is disabled. The Firewall logs (not exactly the most ideal tool but can give insights) are then used to determine what traffic is needed. This is more secure as it takes a deny-all-first approach.

upvoted 1 times

amanshin 1 year, 2 months ago

Option A is a valid approach, but it is not as secure as Option C. In Option A, the application is still exposed to the network, even if it is in an isolated project. This means that if someone were to find a vulnerability in the application, they could potentially exploit it to gain access to the application.

In Option C, the application is isolated from the network by being deployed to a GKE cluster. This means that even if someone were to find a vulnerability in the application, they would not be able to exploit it to gain access to the application.

Additionally, Option C is more scalable and resilient than Option A. This is because a GKE cluster can be scaled up or down as needed, and it is more resistant to failure than a single VM.

Therefore, Option C is the more secure and scalable approach. However, if you are short on time, Option A may be a better option.

upvoted 2 times

Joanale 1 year, 3 months ago

A is a best option, remember you have the hurriest of the contract. Making microservices taking too long and have to know the detailed application architecture. Answer A.

upvoted 2 times

Ric350 1 year, 5 months ago

The answer is A. In real life you would NOT lift and shift an application especially not knowing the ports it uses nor any documentation. That'd be disruptive and cause an outage until you figured it out. You'd be out of a job! The question also clearly states "You want to complete the migratio without putting your environment at risk!"
You'd have to refactor the application in parallel and makes sense if it's a legacy application. You'd want to modernize it with microservices so it can take advantage of all cloud features. If you simply lift and shift, the legacy app cannot take advantage of cloud services so what's the point? You still have the same problems except now you've moved it from on-prem to the cloud.

upvoted 3 times

Ric350 1 year, 5 months ago

Excuse me, C is the correct answer for the reasons listed below. You try lifting and shift a company application without the proper dependenci of how it works, cause a disruption or outage until you figure it out and let me know how that works for you and if you'll still have a job.

upvoted 1 times

sameer2803 1 year, 6 months ago

Answer is B.
even if you disable all traffic within VPC, the request to the application will hit the firewall and will get a deny ingress response. that way we get tc know what port is It coming in. the same can be determined with allowing all traffic in (which exposes your application to the world) but the question ends with "without putting your environment at risk"

upvoted 2 times

pedrojorge 1 year, 7 months ago

Selected Answer: B

B, as A temporarily opens vulnerable paths in the system.

upvoted 3 times

somnathmaddi 1 year, 8 months ago

Selected Answer: A

Answer is A.. You need VPC Flow Logs not "Firewall logs" stated in B

upvoted 4 times

Mixxer5 1 year, 9 months ago

Selected Answer: A

A since B disrupts the system. C and D are out of question if it's supposed to "just work".

upvoted 4 times

Meyucho 1 year, 9 months ago

Selected Answer: B

The difference between A and B is that, in the first, you allow all traffic so the app will work after migration and you can investigate which ports should be open and then take actions. If you go with B you will have a disruption window until figure out all ports needed but will not have any port unneeded port. So... if you asked to avoid disruption go with A and (as in this question) you are asked about security, go with B

upvoted 4 times

pedrojorge 1 year, 7 months ago

The question never asks to avoid disruption, it asks to avoid risk, so the answer must be B.

upvoted 2 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: A

A. Migrate the application into an isolated project using a "Lift & Shift" approach. Enable all internal TCP traffic using VPC Firewall rules. Use VPC Flow logs to determine what traffic should be allowed for the application to work properly.

upvoted 4 times

GPK 2 years, 8 months ago

These questions are no more relevant as google has changed exam and made it really challenging now.

upvoted 1 times

vicky_cyber 2 years, 8 months ago

Could you please help us with recent dumps or guide which dump to be referred

upvoted 2 times

Bwitch 2 years, 7 months ago

This one is accurate.

upvoted 2 times

rr4444 2 years, 8 months ago

Selected Answer: B

B - VPC Flow Logs

Firewall logging only covers TCP and UDP, you explicitly don't know what the app does. That limitation is also important to the fact that implied deny all ingress and deny all egress rules are not covered by Firewall Logging. Plus you have to enable Firewall Logging per rule, so you'd have to have a rule for everything in advance - chicken and egg.... you don't know what is going on, so how could you!?

upvoted 1 times

rr4444 2 years, 8 months ago

VPC FLow logs is A!

I meant A!

upvoted 2 times

keresh 3 years, 3 months ago

Answer A matches "without putting your environment at risk" best, all the other answers are higher in risk

upvoted 4 times

Question #50 *Topic 1*

Your company has deployed an application on Compute Engine. The application is accessible by clients on port 587. You need to balance the load between the different instances running the application. The connection should be secured using TLS, and terminated by the Load Balancer.
What type of Load Balancing should you use?

A. Network Load Balancing

B. HTTP(S) Load Balancing

C. TCP Proxy Load Balancing

D. SSL Proxy Load Balancing

Correct Answer: *D*
Reference:
https://cloud.google.com/load-balancing/docs/ssl/

smart123 Highly Voted 4 years, 1 month ago
Although both TCP Proxy LB and SSL Proxy LB support port 587 but only SSL Proxy LB support TLS. Hence 'D' is the right answer.
upvoted 19 times

umashankar_a Highly Voted 3 years, 1 month ago
Answer D
https://cloud.google.com/load-balancing/docs/ssl
- SSL Proxy Load Balancing is a reverse proxy load balancer that distributes SSL traffic coming from the internet to virtual machine (VM) instances in your Google Cloud VPC network.

When using SSL Proxy Load Balancing for your SSL traffic, user SSL (TLS) connections are terminated at the load balancing layer, and then proxiec to the closest available backend instances by using either SSL (recommended) or TCP.
upvoted 6 times

[Removed] Most Recent 1 year, 1 month ago
Selected Answer: D
"D"
Although port 587 is SMTP (mail) which is an Application Layer protocol, and one might think an Application Layer (HTTPs) Load balancer is needed, according to Google docs, Application Layer LBs offload TLS at GFE which may or may not be the LB. Only the Network Proxy LB confirm: TLS offloading at LB layer. Also, as a general rule, they recommend Network Proxy LB for TLS Offloading:
"..As a general rule, you'd choose an Application Load Balancer when you need a flexible feature set for your applications with HTTP(S) traffic. You choose a proxy Network Load Balancer to implement TLS offload.."

References:
https://cloud.google.com/load-balancing/docs/choosing-load-balancer#flow_chart
https://cloud.google.com/load-balancing/docs/https#control-tls-termination
upvoted 2 times

Ishu_awsguy 1 year, 3 months ago
We can use an HTTPS load balancer and change the backend services port to 587 .|
HTTPS load balacer will also work
upvoted 2 times

Ishu_awsguy 1 year, 3 months ago
accessible by client on port 587 is the power word.
Agree with D
upvoted 1 times

AwesomeGCP 1 year, 11 months ago
Selected Answer: D
Answer D. SSL Proxy Load Balancing
https://cloud.google.com/load-balancing/docs/ssl
upvoted 1 times

dtmtor 3 years, 5 months ago
Answer: D
upvoted 1 times

DebasishLowes 3 years, 6 months ago
Ans : D
upvoted 1 times

[Removed] 3 years, 10 months ago

Ans - D

upvoted 1 times

CHECK666 3 years, 11 months ago

D is the answer. SSL Proxy LoadBalancer supports TLS.

upvoted 2 times

mlyu 4 years ago

Agreed with smart123. Ans is D
https://cloud.google.com/load-balancing/docs/choosing-load-balancer#flow_chart

upvoted 3 times

Question #51 *Topic 1*

You want to limit the images that can be used as the source for boot disks. These images will be stored in a dedicated project.
What should you do?

A. Use the Organization Policy Service to create a compute.trustedimageProjects constraint on the organization level. List the trusted project as the whitelist in an allow operation.

B. Use the Organization Policy Service to create a compute.trustedimageProjects constraint on the organization level. List the trusted projects as the exceptions in a deny operation.

C. In Resource Manager, edit the project permissions for the trusted project. Add the organization as member with the role: Compute Image User.

D. In Resource Manager, edit the organization permissions. Add the project ID as member with the role: Compute Image User.

Correct Answer: *B*
Reference:
https://cloud.google.com/compute/docs/images/restricting-image-access

DebasishLowes Highly Voted 👍 3 years, 5 months ago
Ans : A
upvoted 13 times

[Removed] Highly Voted 👍 3 years, 10 months ago
Ans - A
https://cloud.google.com/compute/docs/images/restricting-image-access#trusted_images
upvoted 8 times

nccdebug Most Recent 🕒 6 months, 2 weeks ago
Correct Answer is: A. Option B suggests listing the trusted projects as exceptions in a deny operation, which is not necessary or recommended. It' simpler and more secure to explicitly allow only the trusted project
upvoted 1 times

Xoxoo 11 months, 2 weeks ago
Selected Answer: A
To limit the images that can be used as the source for boot disks and store these images in a dedicated project, you should use option A:

A. Use the Organization Policy Service to create a compute.trustedimageProjects constraint on the organization level. List the trusted project as th whitelist in an allow operation.

Here's why this option is appropriate:

Organization-Wide Control: Creating an organization-level constraint allows you to enforce the policy organization-wide, ensuring consistent image usage across all projects within the organization.

Whitelist Approach: By listing the trusted project as a whitelist in an "allow" operation, you explicitly specify which project can be trusted as the source for boot disks. This is a more secure approach because it only allows specific trusted projects.

Dedicated Project: You mentioned that the images are stored in a dedicated project, and this option aligns with that requirement.
upvoted 3 times

Xoxoo 11 months, 2 weeks ago
Option B introduces complexity by listing the trusted projects as exceptions in a "deny" operation, which can become challenging to manage more projects are added.
upvoted 1 times

Joanale 1 year, 2 months ago
Actually the default policy is allow * and if you put a constraint it must be as "deny" rule with exceptionsPrincipals or denial conditions. So answer is B, there's no "whitelist".
upvoted 1 times

meh009 1 year, 9 months ago
Selected Answer: A
https://cloud.google.com/compute/docs/images/restricting-image-access#gcloud

Look at the glcoud examples and it will make sense why A is correct
upvoted 3 times

AzureDP900 1 year, 10 months ago

A is right
Use the Trusted image feature to define an organization policy that allows principals to create persistent disks only from images in specific projec

upvoted 2 times

AzureDP900 1 year, 10 months ago

https://cloud.google.com/compute/docs/images/restricting-image-access

upvoted 1 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: A

Answer A. Use the Organization Policy Service to create a compute.trustedimageProjects constraint on the organization level. List the trusted project as the whitelist in an allow operation.

upvoted 2 times

piyush_1982 2 years, 1 month ago

To me the answer seems to be B.
https://cloud.google.com/compute/docs/images/restricting-image-access

By default, instances can be created from images in any project that shares images publicly or explicitly with the user. So there is an implicit allow
Option B states that we need to deny all the projects from being used as a trusted project and add "Trusted Project" as an exception to that rule.

upvoted 4 times

piyush_1982 2 years, 1 month ago

Nope, I think I am getting confused. The correct answer is A.

upvoted 1 times

simbu1299 2 years, 5 months ago

Selected Answer: A

Answer is A

upvoted 2 times

danielklein09 2 years, 5 months ago

Answer is B. You don't whitelist in an allow operation. Since there is an implicit allow, the purpose of the whitelist has been defeated.

upvoted 3 times

gcpengineer 1 year, 3 months ago

implicit deny

upvoted 1 times

CHECK666 3 years, 11 months ago

A is the answer. you need to allow operations.

upvoted 1 times

ownez 4 years ago

I agree with B.

"https://cloud.google.com/compute/docs/images/restricting-image-access"

upvoted 2 times

ownez 4 years ago

Answer is A.

"Use the Trusted image feature to define an organization policy that allows your project members to create persistent disks only from images specific projects."

"After sharing your images with other users, you can control where those users employ those resources within your organization. Set the constraints/compute.storageResourceUseRestrictions constraint to define the projects where users are permitted to use your storage resources."

upvoted 4 times

Sheeda 4 years ago

Yes, A made sense to me too.

upvoted 1 times

Question #52 *Topic 1*

Your team needs to prevent users from creating projects in the organization. Only the DevOps team should be allowed to create projects on behalf of the requester.
Which two tasks should your team perform to handle this request? (Choose two.)

A. Remove all users from the Project Creator role at the organizational level.

B. Create an Organization Policy constraint, and apply it at the organizational level.

C. Grant the Project Editor role at the organizational level to a designated group of users.

D. Add a designated group of users to the Project Creator role at the organizational level.

E. Grant the billing account creator role to the designated DevOps team.

Correct Answer: *BD*

mlyu Highly Voted 4 years ago
I think Ans is AD
Because we need to stop the users can create project first (A), and allow devops team to create project (D)
upvoted 19 times

[Removed] Highly Voted 3 years, 5 months ago
AD is the answer.
If constraint is added , no project creation will be allowed, hence B is wrong
upvoted 7 times

[Removed] Most Recent 1 year, 1 month ago
Selected Answer: AD
"A,D" seems most accurate.
The following page talks about how Project Creator role is granted to all users by default, which is why "A" is necessary. And then there's a sectio
about granting Project Creator to specific users which is where "D" comes in.
https://cloud.google.com/resource-manager/docs/default-access-control#removing-default-roles
upvoted 1 times

AzureDP900 1 year, 10 months ago
AD is perfect.
A. Remove all users from the Project Creator role at the organizational level.
D. Add a designated group of users to the Project Creator role at the organizational level.
upvoted 1 times

AwesomeGCP 1 year, 11 months ago
Selected Answer: AD
A. Remove all users from the Project Creator role at the organizational level.
D. Add a designated group of users to the Project Creator role at the organizational level.

https://cloud.google.com/resource-manager/docs/organization-policy/org-policy-constraints
upvoted 3 times

AzureDP900 1 year, 10 months ago
AD is correct
upvoted 1 times

Jeanphi72 2 years, 1 month ago
Selected Answer: AD
https://cloud.google.com/resource-manager/docs/organization-policy/org-policy-constraints
I see no way to restrict project creation with an organizational policy. If that would have been possible I would have voted for it as restrictions car
be overriden in GCP.
upvoted 4 times

piyush_1982 2 years, 1 month ago
Selected Answer: AC
Seems to be AC
When an organization resource is created, all users in your domain are granted the Billing Account Creator and Project Creator roles by default.
As per the link https://cloud.google.com/resource-manager/docs/default-access-control#removing-default-roles

Hence A is definitely the answer.
Now to add the project creator we need to add the designated group to the project creator role specifically.
upvoted 1 times

absipat 2 years, 2 months ago

ad of course

upvoted 1 times

syllox 3 years, 4 months ago

Ans AC also

upvoted 1 times

> **syllox** 3 years, 4 months ago
>
> AD , C is a mistake it's project Editor and not creator
>
> upvoted 3 times

DebasishLowes 3 years, 6 months ago

Ans : AD

upvoted 4 times

Aniyadu 3 years, 8 months ago

A & D is the right answer.

upvoted 4 times

[Removed] 3 years, 10 months ago

Ans - AD

upvoted 3 times

genesis3k 3 years, 10 months ago

I think AC. Because, a role is granted to user/group, rather user/group is added to a role.

upvoted 1 times

> **syllox** 3 years, 4 months ago
>
> C is a mistake it's project Editor and not creator
>
> upvoted 1 times

CHECK666 3 years, 11 months ago

AD is the answer. There's nothing related to project creation in organization policy constraints.

upvoted 4 times

Question #53 *Topic 1*

A customer deployed an application on Compute Engine that takes advantage of the elastic nature of cloud computing.
How can you work with Infrastructure Operations Engineers to best ensure that Windows Compute Engine VMs are up to date with all the latest OS patches?

A. Build new base images when patches are available, and use a CI/CD pipeline to rebuild VMs, deploying incrementally.

B. Federate a Domain Controller into Compute Engine, and roll out weekly patches via Group Policy Object.

C. Use Deployment Manager to provision updated VMs into new serving Instance Groups (IGs).

D. Reboot all VMs during the weekly maintenance window and allow the StartUp Script to download the latest patches from the internet.

Correct Answer: *D*

genesis3k Highly Voted 3 years, 10 months ago

Answer is A.
Compute Engine doesn't automatically update the OS or the software on your deployed
instances. You will need to patch or update your deployed Compute Engine instances when necessary. However, in the cloud it is not recommended that you patch or update individual running instances. Instead it is best to patch the image that was used to launch the instance ai then replace each affected instance with a new copy.

upvoted 22 times

nccdebug Most Recent 6 months, 2 weeks ago

VM Manager is a suite of tools that can be used to manage operating systems for large virtual machine (VM) fleets running Windows and Linux o Compute Engine.

VM Manager helps drive efficiency through automation and reduces the operational burden of maintaining these VM fleets.

https://cloud.google.com/compute/docs/vm-manager

upvoted 3 times

b6f53d8 8 months ago

Question is outdated, Since 2020 Google has VM Manager for updating VMs (Linux and Windows)

upvoted 2 times

habros 10 months ago

Selected Answer: A

A.

Use a tool like HashiCorp Packer to package the VM images using CI/CD

upvoted 2 times

[Removed] 1 year, 1 month ago

Selected Answer: A

"A"
Applying an OS level patch typically requires a reboot. Rebooting a VM that is actively serving live traffic will have a negative impact on the availability of the service and the user experience and therefore the business.
Out of all the options, only option A emphasises the rolling/gradual deployment of the patch through base images.

References:
https://cloud.google.com/compute/docs/os-patch-management#scheduled_patching

upvoted 2 times

Ric350 1 year, 5 months ago

The answer is definitely D. You would build new base images or deploy new vm's because then you'd have a base OS server with no application o it. You'd have to re-install the app, configure and it as well. You'd have to find a maintenance window that allows you to patch the server, not re-build it! Even the OS patch management doc link below mentions scheduling a time or doing it on demand. You schedule prod systems and patcl the dev/test/staging server on demand bc it's not production. Think practically here. D is the obvious answer.

upvoted 2 times

Ric350 1 year, 5 months ago

correction "would NOT"

upvoted 1 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: A

A. Build new base images when patches are available, and use a CI/CD pipeline to rebuild VMs, deploying incrementally.

upvoted 2 times

PATILDXB 1 year, 8 months ago

you cannot use CI/CD pipeline for building VMs. It is used only for code deployment. Further, building base images is only 1 time activity, organisations cannot afford to change the base image everytime when a patch is released. So, C is the answer

upvoted 1 times

ftpt 1 year ago

you can use CICD with terraform to create new VMs

upvoted 1 times

gcpengineer 1 year, 3 months ago

i use ci/cd to build vm

upvoted 1 times

Aiffone 2 years, 2 months ago

C is obviouly the answer, MIGs help you make sure mahcines deployed are latest image if you want, what's more, its meant to be an elastic syster nothing doesthat better than MIGs.

upvoted 1 times

Jeanphi72 2 years, 1 month ago

Not sure Deployment Manager can indeed create a new MIG and can configure a new deployment of machines with latest OS but what about the existing ones? In addition how to make sure rollout will be smooth?
Option A seems more realistic.

upvoted 2 times

VenkatGCP1 2 years, 8 months ago

The answer is A, we are using this in practice as a solution from Google in one of the top 5 banks for managing windows image patching.

upvoted 4 times

AzureDP900 1 year, 10 months ago

Agreed.

upvoted 1 times

lxs 2 years, 9 months ago

Selected Answer: A

Definitely it will be A. The solution must take the advantage of elasticity of compute engine, so you create a template with patched OS base and redeploy images.

upvoted 2 times

sc_cloud_learn 3 years, 2 months ago

Answer should be A,
C talks about MIG which may not be always needed

upvoted 1 times

DebasishLowes 3 years, 5 months ago

Ans : A

upvoted 2 times

gu9singg 3 years, 5 months ago

this questions still valid for exam?

upvoted 1 times

umashankar_a 3 years, 1 month ago

yeah....even i'm thinking the same, as we got OS Patch Management Service now in GCP for Patching Compute machines as per requirement.
https://cloud.google.com/compute/docs/os-patch-management.
Not really sure on the answer.

upvoted 4 times

DuncanTu 3 years, 1 month ago

Hi

May I know why C is incorrect?

upvoted 1 times

HateMicrosoft 3 years, 5 months ago

The correct anwser is C.
https://cloud.google.com/deployment-manager/docs/reference/latest/deployments/patch

upvoted 1 times

CloudTrip 3 years, 6 months ago

Given the options here Answer D seems practical

upvoted 1 times

singhjoga 3 years, 8 months ago

B seems the only possible answer. Windows patches are configured using Group Policies on the Windows Domain Controller. All other windows machines should be part of the same domain.

upvoted 1 times

FatCharlie 3 years, 9 months ago

The answer is A. This is referring to VMs in an instance group which has built in roll out deployment of new images that can easily be integrated into a CI/CD pipeline.

The people mentioning the patch management tool are considering these to be long running VMs, but that makes little sense in an instance group.

upvoted 3 times

[Removed] 3 years, 10 months ago

Ans - A

upvoted 3 times

Question #54 *Topic 1*

Your team needs to make sure that their backend database can only be accessed by the frontend application and no other instances on the network.
How should your team design this network?

A. Create an ingress firewall rule to allow access only from the application to the database using firewall tags.

B. Create a different subnet for the frontend application and database to ensure network isolation.

C. Create two VPC networks, and connect the two networks using Cloud VPN gateways to ensure network isolation.

D. Create two VPC networks, and connect the two networks using VPC peering to ensure network isolation.

Correct Answer: *A*

singhjoga Highly Voted 3 years, 8 months ago
Although A is correct, but B would be more secure when combined with firewall rules to restrict traffic based on subnets.
Ideal solution would be to use Service Account based firewall rules instead of tag based. See the below paragraph from https://cloud.google.com/solutions/best-practices-vpc-design

"However, even though it is possible to uses tags for target filtering in this manner, we recommend that you use service accounts where possible. Target tags are not access-controlled and can be changed by someone with the instanceAdmin role while VMs are in service. Service accounts are access-controlled, meaning that a specific user must be explicitly authorized to use a service account. There can only be one service account per instance, whereas there can be multiple tags. Also, service accounts assigned to a VM can only be changed when the VM is stopped"
upvoted 7 times

ThisisJohn 2 years, 8 months ago
You may be right but B doesn't mention anything about firewall rules, thus we need to assume there will be communication between both subnets
upvoted 2 times

Aiffone 2 years, 2 months ago
I'm inclined to go with A too because without firewall rules the subnets in B would ensure there is no communication at all due to default implicit rules.
upvoted 1 times

CHECK666 Highly Voted 3 years, 11 months ago
A is the answer, use network tags.
upvoted 6 times

[Removed] Most Recent 1 year, 1 month ago
Selected Answer: A
"A"
The choice is between A and B. Even though subnet isolation is recommended (which would make B correct), subnet isolation alone without accompanying firewall rules does not ensure security.
Only A emphasizes the use of firewall which makes it more correct than B.

Reference:
https://cloud.google.com/architecture/best-practices-vpc-design#target_filtering
upvoted 3 times

Portugapt 5 months, 2 weeks ago
But here the question goes into the design of the network, not the specific implementation details. For design, B makes more sense.
upvoted 1 times

AzureDP900 1 year, 10 months ago
A is correct , rest of the answers doesn't make any sence
upvoted 1 times

azureaspirant 1 year, 9 months ago
@AzureDP900: Cleared AWS Solution Architect Professional (SAP - CO1) on the last date. followed your answers. Cleared 5 GCP Certificates. Glad that you are here.
upvoted 2 times

AwesomeGCP 1 year, 11 months ago
Selected Answer: A
A. Create an ingress firewall rule to allow access only from the application to the database using firewall tags.
upvoted 1 times

zqwiklabs 3 years, 5 months ago

A is definitely incorrect
upvoted 4 times

mistryminded 2 years, 9 months ago
This one is confusing but cannot be A because it says 'Firewall tags'. There is no such thing as firewall tags, only 'Network tags'.
upvoted 2 times

desertlotus1211 3 years, 5 months ago
Answer is D: you'd want the DB in a separate VPC. Allow vpc peering and connect the Front End's backend to the DB. Don't get confused by the question saying 'front end' Front end only means public facing...
upvoted 1 times

Jane111 3 years, 4 months ago
you need to read basic concepts again
upvoted 7 times

AzureDP900 1 year, 10 months ago
A is correct
upvoted 1 times

DebasishLowes 3 years, 6 months ago
Ans : A
upvoted 3 times

[Removed] 3 years, 10 months ago
Ans - A
upvoted 2 times

mlyu 4 years ago
Agree with A
upvoted 2 times

Question #55 *Topic 1*

An organization receives an increasing number of phishing emails.
Which method should be used to protect employee credentials in this situation?

A. Multifactor Authentication

B. A strict password policy

C. Captcha on login pages

D. Encrypted emails

Correct Answer: *D*

DebasishLowes Highly Voted 3 years, 6 months ago

A is the answer.

upvoted 10 times

nccdebug Most Recent 6 months, 2 weeks ago

Ans: A. Implementing MFA helps mitigate the risk posed by phishing attacks by adding an additional barrier to unauthorized access to employee credentials.

upvoted 2 times

[Removed] 1 year, 1 month ago

Selected Answer: A

"A"
Encrypting emails (D) does not prevent or protect against phishing. Phishing leads to attacker getting a user's password. In order to protect again the "impact" of phishing, requiring a second factor would prevent the attacker from logging in using only the password once stolen.

upvoted 3 times

Ric350 1 year, 5 months ago

The question is asking how to PROTECT employees credentials, NOT how to best protect against phishing. MFA does that in case a user's credentials is compromised by have 2FV. It's another defense in layer approach.

upvoted 3 times

Mixxer5 1 year, 9 months ago

Selected Answer: D

MFA itself doesn't really protect user's credentials from beaing leaked. It makes it harder (or nigh impossible) to log in even if they get leaked but they may still leak. Encrypting emails would be of more help, although in case of phishing email it'd be best to educate users and add some filters that will flag external emails as suspicious.

upvoted 1 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: A

A. Multifactor Authentication

upvoted 3 times

GHOST1985 1 year, 11 months ago

Selected Answer: A

https://cloud.google.com/blog/products/g-suite/protecting-you-against-phishing

upvoted 4 times

AzureDP900 1 year, 10 months ago

Agree with A

upvoted 1 times

Deepanshd 1 year, 11 months ago

Selected Answer: A

Multi-factor authentication will prevent employee credentials

upvoted 2 times

fanilgor 1 year, 11 months ago

Selected Answer: A

A for sure

upvoted 1 times

lxs 2 years, 9 months ago

Selected Answer: D

This question has been taken from the GCP book.
upvoted 3 times

mondigo 3 years, 8 months ago

A
https://cloud.google.com/blog/products/g-suite/7-ways-admins-can-help-secure-accounts-against-phishing-g-suite
upvoted 3 times

ronron89 3 years, 8 months ago

https://www.duocircle.com/content/email-security-services/email-security-in-cryptography#:~:text=Customer%20Login-,Email%20Security%20In%20Cryptography%20Is%20One%20Of%20The%20Most,Measures%20To%20event%20Phishing%20Attempts&text=Cybercriminals%20love%20emails%20the%20most,networks%20all%20over%20the%20world.

The answer should be D.
upvoted 2 times

shk2011 3 years, 10 months ago

Logically if i think even if i have not read about cloud answer is A
upvoted 3 times

[Removed] 3 years, 10 months ago

Ans - A
upvoted 2 times

CHECK666 3 years, 11 months ago

The answer is A.
https://cloud.google.com/blog/products/identity-security/protect-users-in-your-apps-with-multi-factor-authentication
upvoted 3 times

Sheeda 4 years ago

Should be A
upvoted 3 times

Question #56 *Topic 1*

A customer is collaborating with another company to build an application on Compute Engine. The customer is building the application tier in their GCP
Organization, and the other company is building the storage tier in a different GCP Organization. This is a 3-tier web application. Communication between portions of the application must not traverse the public internet by any means.
Which connectivity option should be implemented?

A. VPC peering

B. Cloud VPN

C. Cloud Interconnect

D. Shared VPC

Correct Answer: *B*

sc_cloud_learn Highly Voted 👍 3 years, 2 months ago
both are GCP, should be VPC peering- Option A
upvoted 17 times

okhascorpio Most Recent ⊙ 6 months, 2 weeks ago
Selected Answer: C
Key information being "Communication between portions of the application must not traverse the public internet by any means" leaves only option "C" as a valid one, as all other options rely on the public internet for data transmission.
upvoted 1 times

Oujay 2 months ago
Connects your on-premises network to GCP, not relevant for connecting two GCP organizations
upvoted 1 times

[Removed] 8 months, 3 weeks ago
Selected Answer: A
Vpc peering definitely
upvoted 2 times

[Removed] 1 year, 1 month ago
Selected Answer: A
"A"
Since both are in GCP then VPC Peering makes most sense.

References:
https://cloud.google.com/vpc/docs/vpc-peering
upvoted 3 times

shayke 1 year, 10 months ago
Selected Answer: A
only a
upvoted 2 times

AwesomeGCP 1 year, 11 months ago
Selected Answer: A
A – Peering two VPCs does permit traffic to flow between the two shared networks, but it's only bi-directional. Peered VPC networks remain administratively separate.

Dedicated Interconnect connections enable you to connect your on-premises network … in another project, as long as they are both in the same organization. hence A
upvoted 1 times

AzureDP900 1 year, 10 months ago
Agreed, A is correct.
upvoted 1 times

DP_GCP 3 years, 4 months ago
B is not correct because if Cloud VPN is used data travels over internet and question mentions it doesnt want the data to travel through internet. https://cloud.google.com/network-connectivity/docs/vpn/concepts/overview Cloud VPN securely connects your peer network to your Virtual Private Cloud (VPC) network through an IPsec VPN connection. Traffic traveling between the two networks is encrypted by one VPN gateway and then decrypted by the other VPN gateway. This action protects your data as it travels over the internet

upvoted 1 times

PATILDXB 1 year, 8 months ago

Cloud VPN is a private connection, and different from normal IP VPN or IPSecVPN. Cloud VPN does not ride on internet. B is correct and appropriate, as it is cheaper than VPC peering, because VPC peering incurs charges

upvoted 1 times

mikez2023 1 year, 6 months ago

Cloud VPN securely connects your peer network to your Virtual Private Cloud (VPC) network through an IPsec VPN connection. Traffic traveling between the two networks is encrypted by one VPN gateway and then decrypted by the other VPN gateway. This action protects your data as it travels over the internet. You can also connect two instances of Cloud VPN to each other.

upvoted 1 times

nccdebug 6 months, 2 weeks ago

Communication between portions of the application must not traverse the public internet by any means, so A is the answer

upvoted 1 times

dtmtor 3 years, 5 months ago

A, different orgs

upvoted 4 times

DebasishLowes 3 years, 6 months ago

A is the answer.

upvoted 2 times

[Removed] 3 years, 10 months ago

Ans - A

upvoted 3 times

CHECK666 3 years, 11 months ago

A is the ansswer. use VCP Peering.

upvoted 3 times

Akku1614 4 years ago

Yes it Should be VPC Peering. https://cloud.google.com/vpc/docs/vpc-peering

upvoted 3 times

Sheeda 4 years ago

Should be A

upvoted 4 times

Question #57 *Topic 1*

Your team wants to make sure Compute Engine instances running in your production project do not have public IP addresses. The frontend application Compute
Engine instances will require public IPs. The product engineers have the Editor role to modify resources. Your team wants to enforce this requirement.
How should your team meet these requirements?

A. Enable Private Access on the VPC network in the production project.

B. Remove the Editor role and grant the Compute Admin IAM role to the engineers.

C. Set up an organization policy to only permit public IPs for the front-end Compute Engine instances.

D. Set up a VPC network with two subnets: one with public IPs and one without public IPs.

Correct Answer: *C*
Reference:
https://cloud.google.com/compute/docs/ip-addresses/reserve-static-external-ip-address

saurabh1805 Highly Voted 3 years, 10 months ago
C is correct option here, Refer below link for more details.

https://cloud.google.com/resource-manager/docs/organization-policy/org-policy-constraints#constraints-for-specific-services
upvoted 12 times

FatCharlie 3 years, 9 months ago
More specifically, it's the "Restrict VM IP Forwarding" constraint under Compute Engine
upvoted 3 times

FatCharlie 3 years, 9 months ago
Sorry, no. It's the one under that :)

"Define allowed external IPs for VM instances"
upvoted 2 times

AzureDP900 1 year, 10 months ago
Yes, C is right
upvoted 2 times

[Removed] Most Recent 1 year, 1 month ago
Selected Answer: C
"C"
Only C addresses both concerns regarding public IP and the Editor role privileges. Applying constraints at the org level mitigates the editor privileges and provides the access restrictions desired.

References:
https://cloud.google.com/resource-manager/docs/organization-policy/org-policy-constraints#constraints-for-specific-services
upvoted 2 times

passex 1 year, 9 months ago
and how would you want to separate front-end VM's from the other using Org Policy Constraints - IMO option D make more sense
upvoted 4 times

fad3r 1 year, 5 months ago
Intitally I agreed with you but after looking at the link above it does say this.

This list constraint defines the set of Compute Engine VM instances that are allowed to use external IP addresses.
By default, all VM instances are allowed to use external IP addresses.
The allowed/denied list of VM instances must be identified by the VM instance name, in the form: projects/PROJECT_ID/zones/ZONE/instances/INSTANCE

constraints/compute.vmExternalIpAccess

So you can indeed choose with instances have public ips
https://cloud.google.com/resource-manager/docs/organization-policy/org-policy-constraints#constraints-for-specific-services

Define allowed external IPs for VM instances
upvoted 3 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: C

C. Set up an organization policy to only permit public IPs for the front-end Compute Engine instances.

upvoted 4 times

fad3r 1 year, 5 months ago

Intitally I agreed with you but after looking at the link above it does say this.

This list constraint defines the set of Compute Engine VM instances that are allowed to use external IP addresses.
By default, all VM instances are allowed to use external IP addresses.
The allowed/denied list of VM instances must be identified by the VM instance name, in the form:
projects/PROJECT_ID/zones/ZONE/instances/INSTANCE

constraints/compute.vmExternalIpAccess

So you can indeed choose with instances have public ips
https://cloud.google.com/resource-manager/docs/organization-policy/org-policy-constraints#constraints-for-specific-services

Define allowed external IPs for VM instances

upvoted 2 times

fad3r 1 year, 5 months ago

Sorry meant to comment this on the above post

upvoted 1 times

bartlomiejwaw 2 years, 3 months ago

Not C - Editor role is not enough for setting up org policies

upvoted 2 times

DebasishLowes 3 years, 5 months ago

Ans : C

upvoted 3 times

[Removed] 3 years, 10 months ago

Ans - C

upvoted 4 times

HectorLeon2099 3 years, 10 months ago

I'll go with A

upvoted 2 times

Question #58 *Topic 1*

Which two security characteristics are related to the use of VPC peering to connect two VPC networks? (Choose two.)

A. Central management of routes, firewalls, and VPNs for peered networks

B. Non-transitive peered networks; where only directly peered networks can communicate

C. Ability to peer networks that belong to different Google Cloud organizations

D. Firewall rules that can be created with a tag from one peered network to another peered network

E. Ability to share specific subnets across peered networks

Correct Answer: *BC*

DebasishLowes Highly Voted 3 years, 5 months ago

Ans : BC

upvoted 17 times

mlyu Highly Voted 4 years ago

Ans should be BC
https://cloud.google.com/vpc/docs/vpc-peering#key_properties

upvoted 5 times

MohitA 4 years ago

agree BC

upvoted 1 times

ownez 4 years ago

Correct.
B: "Only directly peered networks can communicate. Transitive peering is not supported."

C: " You can make services available privately across different VPC networks within and across organizations."

upvoted 3 times

Mihai89 3 years, 9 months ago

Agree with BC

upvoted 1 times

okhascorpio Most Recent 6 months, 2 weeks ago

Selected Answer: BD

https://cloud.google.com/firewall/docs/tags-firewalls-overview

upvoted 1 times

okhascorpio 6 months, 2 weeks ago

Selected Answer: BD

B and D as the question specifically ask for security capabilities. C is not a security capability while D is.

upvoted 2 times

mackarel22 1 year, 6 months ago

Selected Answer: BC

https://cloud.google.com/vpc/docs/vpc-peering#specifications
Transitive peering is not supported. So BC

upvoted 2 times

Meyucho 1 year, 8 months ago

Selected Answer: CE

Although B is correct, going into detail I think that non-transitivity is just true for networks joined by peering but If there is a third network connected by VPN or Interconnect there is transitivity, so I discard B and stay with C and E

upvoted 1 times

AzureDP900 1 year, 10 months ago

BC is right

upvoted 2 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: BC

B. Non-transitive peered networks; where only directly peered networks can communicate
C. Ability to peer networks that belong to different Google Cloud Platform organizations

upvoted 3 times

zellck 1 year, 11 months ago

Selected Answer: BC

BC is the answer.

upvoted 2 times

Medofree 2 years, 4 months ago

D is false because : "You cannot use a tag or service account from one peered network in the other peered network."

upvoted 1 times

dtmtor 3 years, 5 months ago

Answer is BC

upvoted 3 times

Aniyadu 3 years, 8 months ago

B&C is the right answer

upvoted 2 times

FatCharlie 3 years, 9 months ago

The answers marked in the question seem to be referring to _shared_ VPC capabilities.

upvoted 1 times

[Removed] 3 years, 10 months ago

Ans - BC

upvoted 2 times

CHECK666 3 years, 11 months ago

BC is the answer.

upvoted 2 times

cipher90 4 years ago

AD is correct "Security Characteristics"

upvoted 1 times

mte_tech34 3 years, 11 months ago

No it's not. "You cannot use a tag or service account from one peered network in the other peered network." -> https://cloud.google.com/vpc/docs/vpc-peering

upvoted 2 times

Question #59 *Topic 1*

A patch for a vulnerability has been released, and a DevOps team needs to update their running containers in Google Kubernetes Engine (GKE).
How should the DevOps team accomplish this?

A. Use Puppet or Chef to push out the patch to the running container.

B. Verify that auto upgrade is enabled; if so, Google will upgrade the nodes in a GKE cluster.

C. Update the application code or apply a patch, build a new image, and redeploy it.

D. Configure containers to automatically upgrade when the base image is available in Container Registry.

Correct Answer: *C*
Reference:
https://cloud.google.com/kubernetes-engine/docs/security-bulletins

TNT87 Highly Voted 👍 3 years, 6 months ago
https://cloud.google.com/containers/security
Containers are meant to be immutable, so you deploy a new image in order to make changes. You can simplify patch management by rebuilding your images regularly, so the patch is picked up the next time a container is deployed. Get the full picture of your environment with regular imag
security reviews.
C is better
upvoted 15 times

AzureDP900 1 year, 10 months ago
Yes, C is correct
upvoted 1 times

DebasishLowes Highly Voted 👍 3 years, 5 months ago
Ans : C
upvoted 7 times

GCBC Most Recent 🕓 1 year ago
C is ans - no auto upgrade will patch
upvoted 2 times

[Removed] 1 year, 1 month ago
Selected Answer: C
"C"
Containers are immutable and cannot be updated in place. Base image/container must be patched and then gradually introduced to live containe
pool.

References:
https://cloud.google.com/architecture/best-practices-for-operating-containers#immutability
upvoted 2 times

Ishu_awsguy 1 year, 3 months ago
My vote for B.
This is a biog value add of GKE - inplace upgrades.
upvoted 1 times

Ric350 1 year, 5 months ago
B is 100% the answer.
Fixing some vulnerabilities requires only a control plane upgrade, performed automatically by Google on GKE, while others require both control plane and node upgrades.

To keep clusters patched and hardened against vulnerabilities of all severities, we recommend using node auto-upgrade on GKE (on by default).
https://cloud.google.com/kubernetes-engine/docs/resources/security-patching#how_vulnerabilities_are_patched
upvoted 2 times

AwesomeGCP 1 year, 11 months ago
Selected Answer: C
C. Update the application code or apply a patch, build a new image, and redeploy it.
upvoted 1 times

Medofree 2 years, 4 months ago
Selected Answer: C
Correct ans is C, because "DevOps team needs to update their running containers".
upvoted 2 times

Rhehehe 2 years, 8 months ago

Its actually B.
Patching a vulnerability involves upgrading to a new GKE or Anthos version number. GKE and Anthos versions include versioned components for the operating system, Kubernetes components, and other containers that make up the Anthos platform. Fixing some vulnerabilities requires only control plane upgrade, performed automatically by Google on GKE, while others require both control plane and node upgrades.

To keep clusters patched and hardened against vulnerabilities of all severities, we recommend using node auto-upgrade on GKE (on by default). C
other Anthos platforms, Google recommends upgrading your Anthos components at least monthly.

Ref: https://cloud.google.com/kubernetes-engine/docs/resources/security-patching

upvoted 5 times

StanPeng 2 years, 6 months ago

The qeustion is asking about upgrading application code rather than GKE

upvoted 1 times

Ric350 1 year, 5 months ago

No, the question is asking how vulnerabilities are patched! To keep clusters patched and hardened against vulnerabilities of all severities, w
recommend using node auto-upgrade on GKE (on by default).
https://cloud.google.com/kubernetes-engine/docs/resources/security-patching#how_vulnerabilities_are_patched

upvoted 2 times

alexm112 2 years, 6 months ago

Agreed - I think this wasn't available at the time people responded.

B is correct
https://cloud.google.com/kubernetes-engine/docs/how-to/node-auto-upgrades

upvoted 2 times

SuperDevops 2 years, 9 months ago

I took the test yesterday and didn't pass, NO ISSUE is from here. The questions are totally new
Whizlabs it´s OK

upvoted 1 times

sriz 2 years, 9 months ago

u got questions from Whizlabs?

upvoted 2 times

Aniyadu 3 years, 8 months ago

The question asked is "team needs to update their running containers" if its was auto enabled there was no need to update manually. so my answer will be C.

upvoted 2 times

Kevinsayn 3 years, 9 months ago

Me voy definitivamente con la C, dado que actualizar los nodos con autoupgrade no tiene nada que ver con los contenedores, la vulnerabilidad e
este caso se debe aplicar con respecto a contenedor ósea aplicación por lo que la respuesta C es la correcta.

upvoted 3 times

soukumar369 3 years, 9 months ago

Translaed : 'm definitely going with C, since updating the nodes with autoupgrade has nothing to do with the containers, the vulnerability in th
case must be applied with respect to the application bone container so the C answer is correct.

upvoted 1 times

jonclem 3 years, 9 months ago

Answer B is correct as per the Video Google Kubernetes Engine (GKE) Security on Linuxacademy.

upvoted 2 times

[Removed] 3 years, 10 months ago

Ans - C

upvoted 3 times

Rantu 3 years, 11 months ago

C is the correct answer as this is the way to patch, build, re-deploy

upvoted 3 times

Namaste 3 years, 11 months ago

Answer is C.

upvoted 3 times

ownez 3 years, 11 months ago

I would go for C because some reported CVEs will take time to be published and approval in CVE advisory portal. Once approved, it will notify to necessary third party.

Hence, this requires a lot of time and left people exposed to the vulnerability.

Answer is C.

upvoted 3 times

Question #60 *Topic 1*

A company is running their webshop on Google Kubernetes Engine and wants to analyze customer transactions in BigQuery. You need to ensure that no credit card numbers are stored in BigQuery
What should you do?

A. Create a BigQuery view with regular expressions matching credit card numbers to query and delete affected rows.

B. Use the Cloud Data Loss Prevention API to redact related infoTypes before data is ingested into BigQuery.

C. Leverage Security Command Center to scan for the assets of type Credit Card Number in BigQuery.

D. Enable Cloud Identity-Aware Proxy to filter out credit card numbers before storing the logs in BigQuery.

Correct Answer: *D*

saurabh1805 Highly Voted 👍 3 years, 10 months ago
B is correct answer here.
upvoted 12 times

saurabh1805 3 years, 10 months ago
https://cloud.google.com/bigquery/docs/scan-with-dlp
upvoted 4 times

jhkkrishnan Most Recent ⊙ 1 month ago
sdfdfwerrwerweewrwr
upvoted 1 times

pixfw1 2 months, 2 weeks ago
DLP for sure.
upvoted 1 times

madcloud32 4 months, 3 weeks ago
Selected Answer: B
B is correct.
got this in exam. Dump is valid. Few new came but easy ones.
upvoted 1 times

cloud_monk 5 months, 1 week ago
Selected Answer: B
DLP is the service specifically for this task.
upvoted 1 times

madcloud32 6 months ago
Selected Answer: B
B is correct. DLP
upvoted 1 times

[Removed] 8 months, 3 weeks ago
Selected Answer: B
B - you want to use dlp for that
upvoted 2 times

jsiror 1 year ago
Selected Answer: B
B is the correct answer
upvoted 2 times

[Removed] 1 year, 1 month ago
Selected Answer: B
"B"
A and C are reactive measures. D is not related to hiding sensitive information. B is the only pro-active/preventative measure specific to hiding sensitive information.

https://cloud.google.com/bigquery/docs/scan-with-dlp
upvoted 2 times

pedrojorge 1 year, 7 months ago
Selected Answer: B

B.
https://cloud.google.com/bigquery/docs/scan-with-dlp

upvoted 2 times

jaykumarjkd99 1 year, 8 months ago

Selected Answer: B

B is correct answer here.
.

upvoted 2 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: B

B. Use the Cloud Data Loss Prevention API to redact related infoTypes before data is ingested into BigQuery.

upvoted 3 times

giovy_82 2 years ago

Selected Answer: B

How can it be D? i'll go for B, DLP is the tool to scan and find sensible data

upvoted 1 times

sudarchary 2 years, 7 months ago

https://cloud.google.com/bigquery/docs/scan-with-dlp

upvoted 1 times

sudarchary 2 years, 7 months ago

Selected Answer: B

Cloud Data Loss Prevention API allows to detect and redact or remove
sensitive data before the comments or reviews are published. Cloud DLP will read
information from BigQuery, Cloud Storage or Datastore and scan it for sensitive data.

upvoted 1 times

AzureDP900 1 year, 10 months ago

B is correct

upvoted 1 times

rr4444 2 years, 8 months ago

Selected Answer: B

D is silly

upvoted 1 times

[Removed] 3 years, 4 months ago

D is impossible. I support B

upvoted 2 times

Question #61 *Topic 1*

A customer wants to deploy a large number of 3-tier web applications on Compute Engine.
How should the customer ensure authenticated network separation between the different tiers of the application?

A. Run each tier in its own Project, and segregate using Project labels.

B. Run each tier with a different Service Account (SA), and use SA-based firewall rules.

C. Run each tier in its own subnet, and use subnet-based firewall rules.

D. Run each tier with its own VM tags, and use tag-based firewall rules.

Correct Answer: *C*

genesis3k Highly Voted 3 years, 10 months ago

Answer is B. Keyword is 'authenticated". Reference below:
"Isolate VMs using service accounts when possible"
"even though it is possible to uses tags for target filtering in this manner, we recommend that you use service accounts where possible. Target tag are not access-controlled and can be changed by someone with the instanceAdmin role while VMs are in service. Service accounts are access-controlled, meaning that a specific user must be explicitly authorized to use a service account. There can only be one service account per instance whereas there can be multiple tags. Also, service accounts assigned to a VM can only be changed when the VM is stopped."
https://cloud.google.com/solutions/best-practices-vpc-design#isolate-vms-service-accounts

upvoted 30 times

Ric350 1 year, 5 months ago

Thank you for this great explanation with link to documentation.

upvoted 1 times

gu9singg 3 years, 5 months ago

document says about subnet isolation

upvoted 2 times

AzureDP900 1 year, 10 months ago

Agreed with you and B is right

upvoted 1 times

pico Most Recent 3 months, 3 weeks ago

Selected Answer: C

why the other options are less ideal:

A. Project labels: Project labels are primarily for organizational purposes and don't provide strong network isolation.
B. Service Accounts: While service accounts can be used for authentication, using them alone for network separation can be complex and less effective than subnet-based rules.
D. VM tags: VM tags can be used for filtering in firewall rules, but they don't inherently create network separation.

upvoted 1 times

ArizonaClassics 11 months, 3 weeks ago

Run each tier with a different Service Account (SA), and use SA-based firewall rules: Service accounts are primarily designed for authentication and authorization of service-to-service interactions. Using them for network separation is possible but is not their primary use case.

D. Run each tier with its own VM tags, and use tag-based firewall rules: This is the most recommended method for multi-tier applications. VM tag are a straightforward way to identify the role or purpose of a VM (like 'web', 'app', 'database'). When VMs are tagged appropriately, tag-based firewall rules can easily control which tiers can communicate with each other. For example, firewall rules can be set so that only VMs with the 'web tag can communicate with VMs with the 'app' tag, and so on.

upvoted 2 times

GCBC 1 year ago

B - https://cloud.google.com/solutions/best-practices-vpc-design#isolate-vms-service-accounts

upvoted 2 times

[Removed] 1 year, 1 month ago

Selected Answer: B

"B"
Keyword here is "authenticated". Service account related answer is the only option that addresses authentication. The rest are network security related.

References:
https://cloud.google.com/compute/docs/access/service-accounts#use-sas
https://cloud.google.com/solutions/best-practices-vpc-design#isolate-vms-service-accounts

upvoted 4 times

riteshahir5815 1 year, 5 months ago

Selected Answer: C

c is correct answer.

upvoted 2 times

mahi9 1 year, 6 months ago

Selected Answer: B

SA accounts

upvoted 1 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: B

B. Run each tier with a different Service Account (SA), and use SA-based firewall rules.

upvoted 1 times

mynk29 2 years, 6 months ago

"As previously mentioned, you can identify the VMs on a specific subnet by applying a unique network tag or service account to those instances. This allows you to create firewall rules that only apply to the VMs in a subnet—those with the associated network tag or service account. For example, to create a firewall rule that permits all communication between VMs in the same subnet, you can use the following rule configuration c the Firewall rules page:"

B is the right answer

upvoted 2 times

mistryminded 2 years, 9 months ago

Selected Answer: B

Answer is B - https://cloud.google.com/vpc/docs/firewalls#service-accounts-vs-tags

upvoted 2 times

gu9singg 3 years, 5 months ago

C: is incorrect, we need to authenticate, network rules does not apply and not a recommend best practice from google

upvoted 2 times

> **gu9singg** 3 years, 5 months ago
>
> C: is incorrect because we need to spend lot of time designing the network topology etc, google recommended practice is to use simple network design with automation in mind, so service account provides those, hence final decision goes to B
>
> upvoted 2 times

> **gu9singg** 3 years, 5 months ago
>
> Correct answer is B
>
> upvoted 2 times

DebasishLowes 3 years, 5 months ago

Ans : C

upvoted 2 times

singhjoga 3 years, 8 months ago

B as per best practices https://cloud.google.com/solutions/best-practices-vpc-design

upvoted 3 times

Fellipo 3 years, 9 months ago

B exists?

upvoted 1 times

[Removed] 3 years, 10 months ago

Ans - C
https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations#networking_and_security

https://cloud.google.com/solutions/best-practices-vpc-design#addresses_and_subnets

upvoted 2 times

Rantu 3 years, 11 months ago

Authenticated separation is the key here. Ideally it should be tag based firewall rule separation. However authenticated word creates confusion. M best judgement is B

upvoted 2 times

CHECK666 3 years, 11 months ago

C is the answer.

upvoted 2 times

Question #62 *Topic 1*

A manager wants to start retaining security event logs for 2 years while minimizing costs. You write a filter to select the appropriate log entries.
Where should you export the logs?

A. BigQuery datasets

B. Cloud Storage buckets

C. StackDriver logging

D. Cloud Pub/Sub topics

Correct Answer: *C*
Reference:
https://cloud.google.com/logging/docs/exclusions

madcloud32 6 months ago

Selected Answer: B

B : GCS without any doubts.

upvoted 2 times

[Removed] 8 months, 3 weeks ago

Selected Answer: B

B - minimizing cost

upvoted 3 times

[Removed] 1 year, 1 month ago

Selected Answer: B

"B"
Keyword here is minimizing cost. Cloud storage is typically the most cost effective option.

References:
https://cloud.google.com/blog/products/storage-data-transfer/how-to-save-on-google-cloud-storage-costs

upvoted 3 times

shayke 1 year, 8 months ago

Selected Answer: B

B- is the cheapest optaion

upvoted 2 times

AzureDP900 1 year, 10 months ago

B is best for cost optimization perspective

upvoted 2 times

shayke 1 year, 10 months ago

Selected Answer: B

GCS would be the chipest option

upvoted 2 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: B

B. Cloud Storage buckets

upvoted 1 times

Deepanshd 1 year, 11 months ago

Selected Answer: B

Cloud storage is always considered when minimize cost

upvoted 1 times

Bill1000 1 year, 11 months ago

B is correct

upvoted 2 times

mbiy 2 years, 6 months ago

Ans C is correct, you can define a custom log bucket and mention the retention policy for any number of years (range - 1 day to 3650 days). Underlying these custom define log bucket is also created within Cloud Storage. As per the question you can retain log for 2 years in Stackdriver Logging which is aka Cloud Logging, and then later archive to cold line storage if there is a requirement.

upvoted 1 times

VJ_0909 2 years, 6 months ago

Default retention for logging is 30 days because it is expensive to hold the logs there for longer duration. Bucket is always the cheapest option

upvoted 1 times

jayk22 2 years, 10 months ago

Ans B. Validated.

upvoted 4 times

DebasishLowes 3 years, 5 months ago

Ans: B

upvoted 4 times

[Removed] 3 years, 10 months ago

Ans - B

upvoted 1 times

Raushanr 3 years, 11 months ago

Ans is B

upvoted 1 times

mlyu 4 years ago

Ans B
Cloud storage is always considered when minimize cost

upvoted 2 times

MohitA 4 years ago

Agree B

upvoted 1 times

Question #63 *Topic 1*

For compliance reasons, an organization needs to ensure that in-scope PCI Kubernetes Pods reside on `in-scope` Nodes only. These Nodes can only contain the
`in-scope` Pods.
How should the organization achieve this objective?

A. Add a nodeSelector field to the pod configuration to only use the Nodes labeled inscope: true.

B. Create a node pool with the label inscope: true and a Pod Security Policy that only allows the Pods to run on Nodes with that label.

C. Place a taint on the Nodes with the label inscope: true and effect NoSchedule and a toleration to match in the Pod configuration.

D. Run all in-scope Pods in the namespace ג€in-scope-pciג€.

Correct Answer: *C*

Tabayashi Highly Voted 2 years, 4 months ago

[A] Correct answer. This is a typical use case for node selector.
https://kubernetes.io/docs/concepts/scheduling-eviction/assign-pod-node/#nodeselector

[B] The Pod Security Policy is designed to block the creation of misconfigured pods on certain clusters. This does not meet the requirements.

[C] Taint will no longer place pods without the "inscope" label on that node, but it does not guarantee that pods with the "inscope" label will be placed on that node.

[D] Placing the "in scope" node in the namespace "in-scope-pci" may meet the requirement, but [A] takes precedence.

upvoted 11 times

AzureDP900 1 year, 10 months ago

A is correct.

upvoted 1 times

gcpengineer 1 year, 3 months ago

C is correct

upvoted 3 times

MariaGabiGabriela 2 years, 3 months ago

I think [A] does not stop other pods from being run in the PCI node, which is a requirement as the question states... I would go with [C]

upvoted 7 times

gcpengineer Highly Voted 1 year, 3 months ago

Selected Answer: C

C is the ans as per chatgpt

upvoted 6 times

pico Most Recent 3 months, 2 weeks ago

Selected Answer: C

why the other options are less suitable:

A. nodeSelector: While nodeSelector can help target pods to specific nodes, it doesn't prevent other pods from being scheduled on those nodes they fit the node's resources.
B. Node pool and Pod Security Policy: Pod Security Policies are deprecated in newer Kubernetes versions, and node pools alone won't guarantee the required isolation.
D. Namespace: Namespaces provide logical separation but don't inherently enforce node-level restrictions.

upvoted 1 times

rsamant 9 months ago

A
https://kubernetes.io/docs/concepts/scheduling-eviction/assign-pod-node/

upvoted 1 times

ArizonaClassics 11 months, 3 weeks ago

C. Place a taint on the Nodes with the label inscope: true and effect NoSchedule and a toleration to match in the Pod configuration: This is the be solution. Taints and tolerations work together to ensure that Pods are not scheduled onto inappropriate nodes. By placing a taint on the Nodes, you are essentially marking them so that they repel all Pods that don't have a matching toleration. With this method, only Pods with the correct toleration can be scheduled on in-scope Nodes, ensuring compliance.

upvoted 2 times

Meyucho 1 year, 8 months ago

Selected Answer: C

A nodeselector configuration is from a pod template perspective. This question ask to PRESERVE some nodes for specific pods, so this is the main utilization for TAINT. This is a conceptual question and the answer is C

upvoted 4 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: A

A. Add a nodeSelector field to the pod configuration to only use the Nodes labeled inscope: true.

upvoted 3 times

GHOST1985 1 year, 11 months ago

Selected Answer: A

nodeSelector is the simplest recommended form of node selection constraint. You can add the nodeSelector field to your Pod specification and specify the node labels you want the target node to have. Kubernetes only schedules the Pod onto nodes that have each of the labels you specify => https://kubernetes.io/docs/concepts/scheduling-eviction/assign-pod-node/#nodeselector

Tolerations are applied to pods. Tolerations allow the scheduler to schedule pods with matching taints. Tolerations allow scheduling but don't guarantee scheduling: the scheduler also evaluates other parameters as part of its function. => https://kubernetes.io/docs/concepts/scheduling-eviction/taint-and-toleration/

upvoted 3 times

fanilgor 1 year, 11 months ago

Selected Answer: C

Basic K8s principles of scheduling workloads.
Taints and tolerations make perfect sense for this use case. Therefore C.

upvoted 2 times

Jeanphi72 2 years ago

Selected Answer: A

https://redhat-scholars.github.io/kubernetes-tutorial/kubernetes-tutorial/taints-affinity.html
A Taint is applied to a Kubernetes Node that signals the scheduler to avoid or not schedule certain Pods.
A Toleration is applied to a Pod definition and provides an exception to the taint.

https://kubernetes.io/docs/concepts/scheduling-eviction/taint-and-toleration/
Node affinity is a property of Pods that attracts them to a set of nodes (either as a preference or a **hard requirement**).
Taints are the opposite -- they allow a node to repel a set of pods.

upvoted 3 times

hybridpro 2 years, 2 months ago

Answer should be C. "These Nodes can only contain the
â€in-scopeâ€ Pods." - this can only be achieved by taints and tolerations.

upvoted 1 times

Question #64 *Topic 1*

In an effort for your company messaging app to comply with FIPS 140-2, a decision was made to use GCP compute and network services. The messaging app architecture includes a Managed Instance Group (MIG) that controls a cluster of Compute Engine instances. The instances use Local SSDs for data caching and
UDP for instance-to-instance communications. The app development team is willing to make any changes necessary to comply with the standard
Which options should you recommend to meet the requirements?

A. Encrypt all cache storage and VM-to-VM communication using the BoringCrypto module.

B. Set Disk Encryption on the Instance Template used by the MIG to customer-managed key and use BoringSSL for all data transit between instances.

C. Change the app instance-to-instance communications from UDP to TCP and enable BoringSSL on clients' TLS connections.

D. Set Disk Encryption on the Instance Template used by the MIG to Google-managed Key and use BoringSSL library on all instance-to-instance communications.

Correct Answer: *D*

subhala Highly Voted 3 years, 9 months ago

when I revisited this, Now I think A is correct. In A - We will use an approved encryption method for encrypting Local SSD and VM to VM communication. In B and D, we are still using GCP's encryption algorithms and are not FIPS 140-2 approved. Moreover only the BoringCrypto is FIPS 140-2 approved and not the Boring SSL. I see A as evidently correct. ownez, genesis3k, MohitA has explained this and provided the right link too.

upvoted 14 times

LaithTech Most Recent 3 weeks, 6 days ago

Selected Answer: B

The correct answer is B

upvoted 1 times

3d9563b 1 month, 2 weeks ago

Selected Answer: B

A. Encrypt all cache storage and VM-to-VM communication using the BoringCrypto module:

BoringCrypto is not an established or widely recognized cryptographic library for FIPS 140-2 compliance. Instead, BoringSSL or OpenSSL with FIPS validation should be used for both data-at-rest and data-in-transit encryption.
C. Change the app instance-to-instance communications from UDP to TCP and enable BoringSSL on clients' TLS connections:

While changing from UDP to TCP might provide more reliable connections, it does not directly address FIPS 140-2 compliance. You still need to ensure that all data-in-transit encryption uses a validated cryptographic module such as BoringSSL.
D. Set Disk Encryption on the Instance Template used by the MIG to Google-managed Key and use BoringSSL library on all instance-to-instance communications:

Google-managed keys for disk encryption do not provide the level of control required for FIPS 140-2 compliance, which typically requires customer-managed keys for greater control and accountability.

upvoted 1 times

gical 8 months, 2 weeks ago

Selected answer B
https://cloud.google.com/security/compliance/fips-140-2-validated/
"Google's Local SSD storage product is automatically encrypted with NIST approved ciphers, but Google's current implementation for this produc doesn't have a FIPS 140-2 validation certificate. If you require FIPS-validated encryption on Local SSD storage, you must provide your own encryption with a FIPS-validated cryptographic module."

upvoted 4 times

b6f53d8 8 months ago

YES, as in your link: you need to encrypt SSD using your own solution, and BoringSSL is a library to use

upvoted 1 times

ArizonaClassics 11 months, 3 weeks ago

A. Encrypt all cache storage and VM-to-VM communication using the BoringCrypto module.

This option ensures both storage (Local SSDs) and inter-instance communications are encrypted using a FIPS 140-2 compliant module.

upvoted 4 times

ArizonaClassics 11 months, 3 weeks ago

A. Encrypt all cache storage and VM-to-VM communication using the BoringCrypto module.

This option ensures both storage (Local SSDs) and inter-instance communications are encrypted using a FIPS 140-2 compliant module.

upvoted 1 times

ymkk 1 year ago

Selected Answer: A

https://cloud.google.com/security/compliance/fips-140-2-validated/

upvoted 2 times

gcpengineer 1 year, 3 months ago

Selected Answer: A

A is the ans

upvoted 2 times

pedrojorge 1 year, 7 months ago

Selected Answer: C

"BoringSSL as a whole is not FIPS validated. However, there is a core library (called BoringCrypto) that has been FIPS validated"
https://boringssl.googlesource.com/boringssl/+/master/crypto/fipsmodule/FIPS.md

upvoted 3 times

AzureDP900 1 year, 10 months ago

https://cloud.google.com/docs/security/key-management-deep-dive

A is right

upvoted 1 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: A

A. Encrypt all cache storage and VM-to-VM communication using the BoringCrypto module.

upvoted 1 times

sudarchary 2 years, 7 months ago

Selected Answer: A

FIPS140 module is supported

upvoted 2 times

[Removed] 3 years, 4 months ago

D is the correct answer

upvoted 2 times

DebasishLowes 3 years, 5 months ago

Ans : A

upvoted 1 times

TNT87 3 years, 5 months ago

https://cloud.google.com/security/compliance/fips-140-2-validated
Google Cloud Platform uses a FIPS 140-2 validated encryption module called BoringCrypto (certificate 3318) in our production environment. This means that both data in transit to the customer and between data centers, and data at rest are encrypted using FIPS 140-2 validated encryption. The module that achieved FIPS 140-2 validation is part of our BoringSSL library.
Ans A

upvoted 4 times

TNT87 3 years, 6 months ago

A is the answer https://boringssl.googlesource.com/boringssl/+/master/crypto/fipsmodule/FIPS.md

upvoted 2 times

chetz12 3 years, 8 months ago

I think A is correct as that's the only one support FIPS140 module

upvoted 3 times

Question #65 *Topic 1*

A customer has an analytics workload running on Compute Engine that should have limited internet access.
Your team created an egress firewall rule to deny (priority 1000) all traffic to the internet.
The Compute Engine instances now need to reach out to the public repository to get security updates.
What should your team do?

A. Create an egress firewall rule to allow traffic to the CIDR range of the repository with a priority greater than 1000.

B. Create an egress firewall rule to allow traffic to the CIDR range of the repository with a priority less than 1000.

C. Create an egress firewall rule to allow traffic to the hostname of the repository with a priority greater than 1000.

D. Create an egress firewall rule to allow traffic to the hostname of the repository with a priority less than 1000.

Correct Answer: *C*

dtmtor Highly Voted 👍 3 years, 5 months ago
Answer is B. Lower number is higher priority and dest is only IP ranges in firewall rules
upvoted 25 times

[Removed] Highly Voted 👍 8 months, 3 weeks ago
Selected Answer: B
B... no hostname in firewall rules and lower number = higher priority.
upvoted 5 times

madcloud32 Most Recent 🕒 6 months ago
Selected Answer: B
B is correct.
upvoted 1 times

shayke 1 year, 8 months ago
Selected Answer: B
Ans in B lower number higher priority
upvoted 3 times

Littleivy 1 year, 9 months ago
Selected Answer: B
Answer is B
upvoted 3 times

GHOST1985 1 year, 10 months ago
Selected Answer: B
https://cloud.google.com/vpc/docs/firewalls#priority_order_for_firewall_rules
upvoted 4 times

AzureDP900 1 year, 10 months ago
B is correct
upvoted 2 times

Premumar 1 year, 10 months ago
Selected Answer: B
First filter is priority should be less than 1000. So, option A and C are rejected. Then, we use CIDR range to allow firewall. So, the final answer is B.
upvoted 3 times

AwesomeGCP 1 year, 11 months ago
Selected Answer: B
B. Create an egress firewall rule to allow traffic to the CIDR range of the repository with a priority less than 1000.
Firewall rules only support IPv4 connections. When specifying a source for an ingress rule or a destination for an egress rule by address, you can only use an IPv4 address or IPv4 block in CIDR notation. So Answer is B
upvoted 4 times

piyush_1982 2 years, 1 month ago
Selected Answer: A
The correct answer is A.
As per the link https://cloud.google.com/vpc/docs/firewalls#rule_assignment

Lowest priority in the firewall rule is 65535. So in order for a rule to be of higher priority than 1000 the rule should have a priority of number less than 1000.

upvoted 2 times

Premumar 1 year, 10 months ago

Your explanation is correct. But, option you selected is wrong. It has to be option B.

upvoted 3 times

Rithac 3 years, 2 months ago

I think I am confusing myself by overthinking the wording of this question. I know the answer is A or B since "using hostname is not one of the options in firewall egress rule destination" I also know that "The firewall rule priority is an integer from 0 to 65535, inclusive. Lower integers indica higher priorities." I know that I could resolve this by setting TCP port 80 rule to a priority of 500 (smaller number, but higher priority) and be done Where i'm second guessing myself, is Google referring to the integer or strictly priority? If integer then i'd choose B "priority less than 1000 (small number)", if priority then i'd choose A "priority greater than 1000" (still the lower number). Have I thoroughly confused this question? I"m leaning toward the answer being "A:

upvoted 5 times

DebasishLowes 3 years, 5 months ago

Ans : B

upvoted 3 times

ronron89 3 years, 8 months ago

Answer: B
https://cloud.google.com/vpc/docs/firewalls#rule_assignment
The priority of the second rule determines whether TCP traffic to port 80 is allowed for the webserver targets:

If the priority of the second rule is set to a number greater than 1000, it has a lower priority, so the first rule denying all traffic applies.

If the priority of the second rule is set to 1000, the two rules have identical priorities, so the first rule denying all traffic applies.

If the priority of the second rule is set to a number less than 1000, it has a higher priority, thus allowing traffic on TCP 80 for the webserver target: Absent other rules, the first rule would still deny other types of traffic to the webserver targets, and it would also deny all traffic, including TCP 80, to instances without the webserver tag.

upvoted 4 times

[Removed] 3 years, 10 months ago

Ans - B

upvoted 3 times

Raushanr 3 years, 11 months ago

The firewall rule priority is an integer from 0 to 65535, inclusive. Lower integers indicate higher priorities. If you do not specify a priority when creating a rule, it is assigned a priority of 1000.

upvoted 1 times

Raushanr 3 years, 11 months ago

Answer-B

upvoted 4 times

ownez 4 years ago

It should be B.

Firewall rules can be only specify by IPv4 address or IPv4 block in CIDR.
And it must be lesser priority than 1000 because if more than that, it will overwrite the deny rule.

upvoted 2 times

ownez 3 years, 12 months ago

Sorry It should be A.

"Priority: the numeric evaluation order of the rule. A rule with a priority of 1 is evaluated first. Priorities must be unique for each rule. A good practice is to give rules priority numbers that allow later insertion (such as 100, 200, 300)."

upvoted 1 times

ownez 3 years, 11 months ago

Correction. B

upvoted 3 times

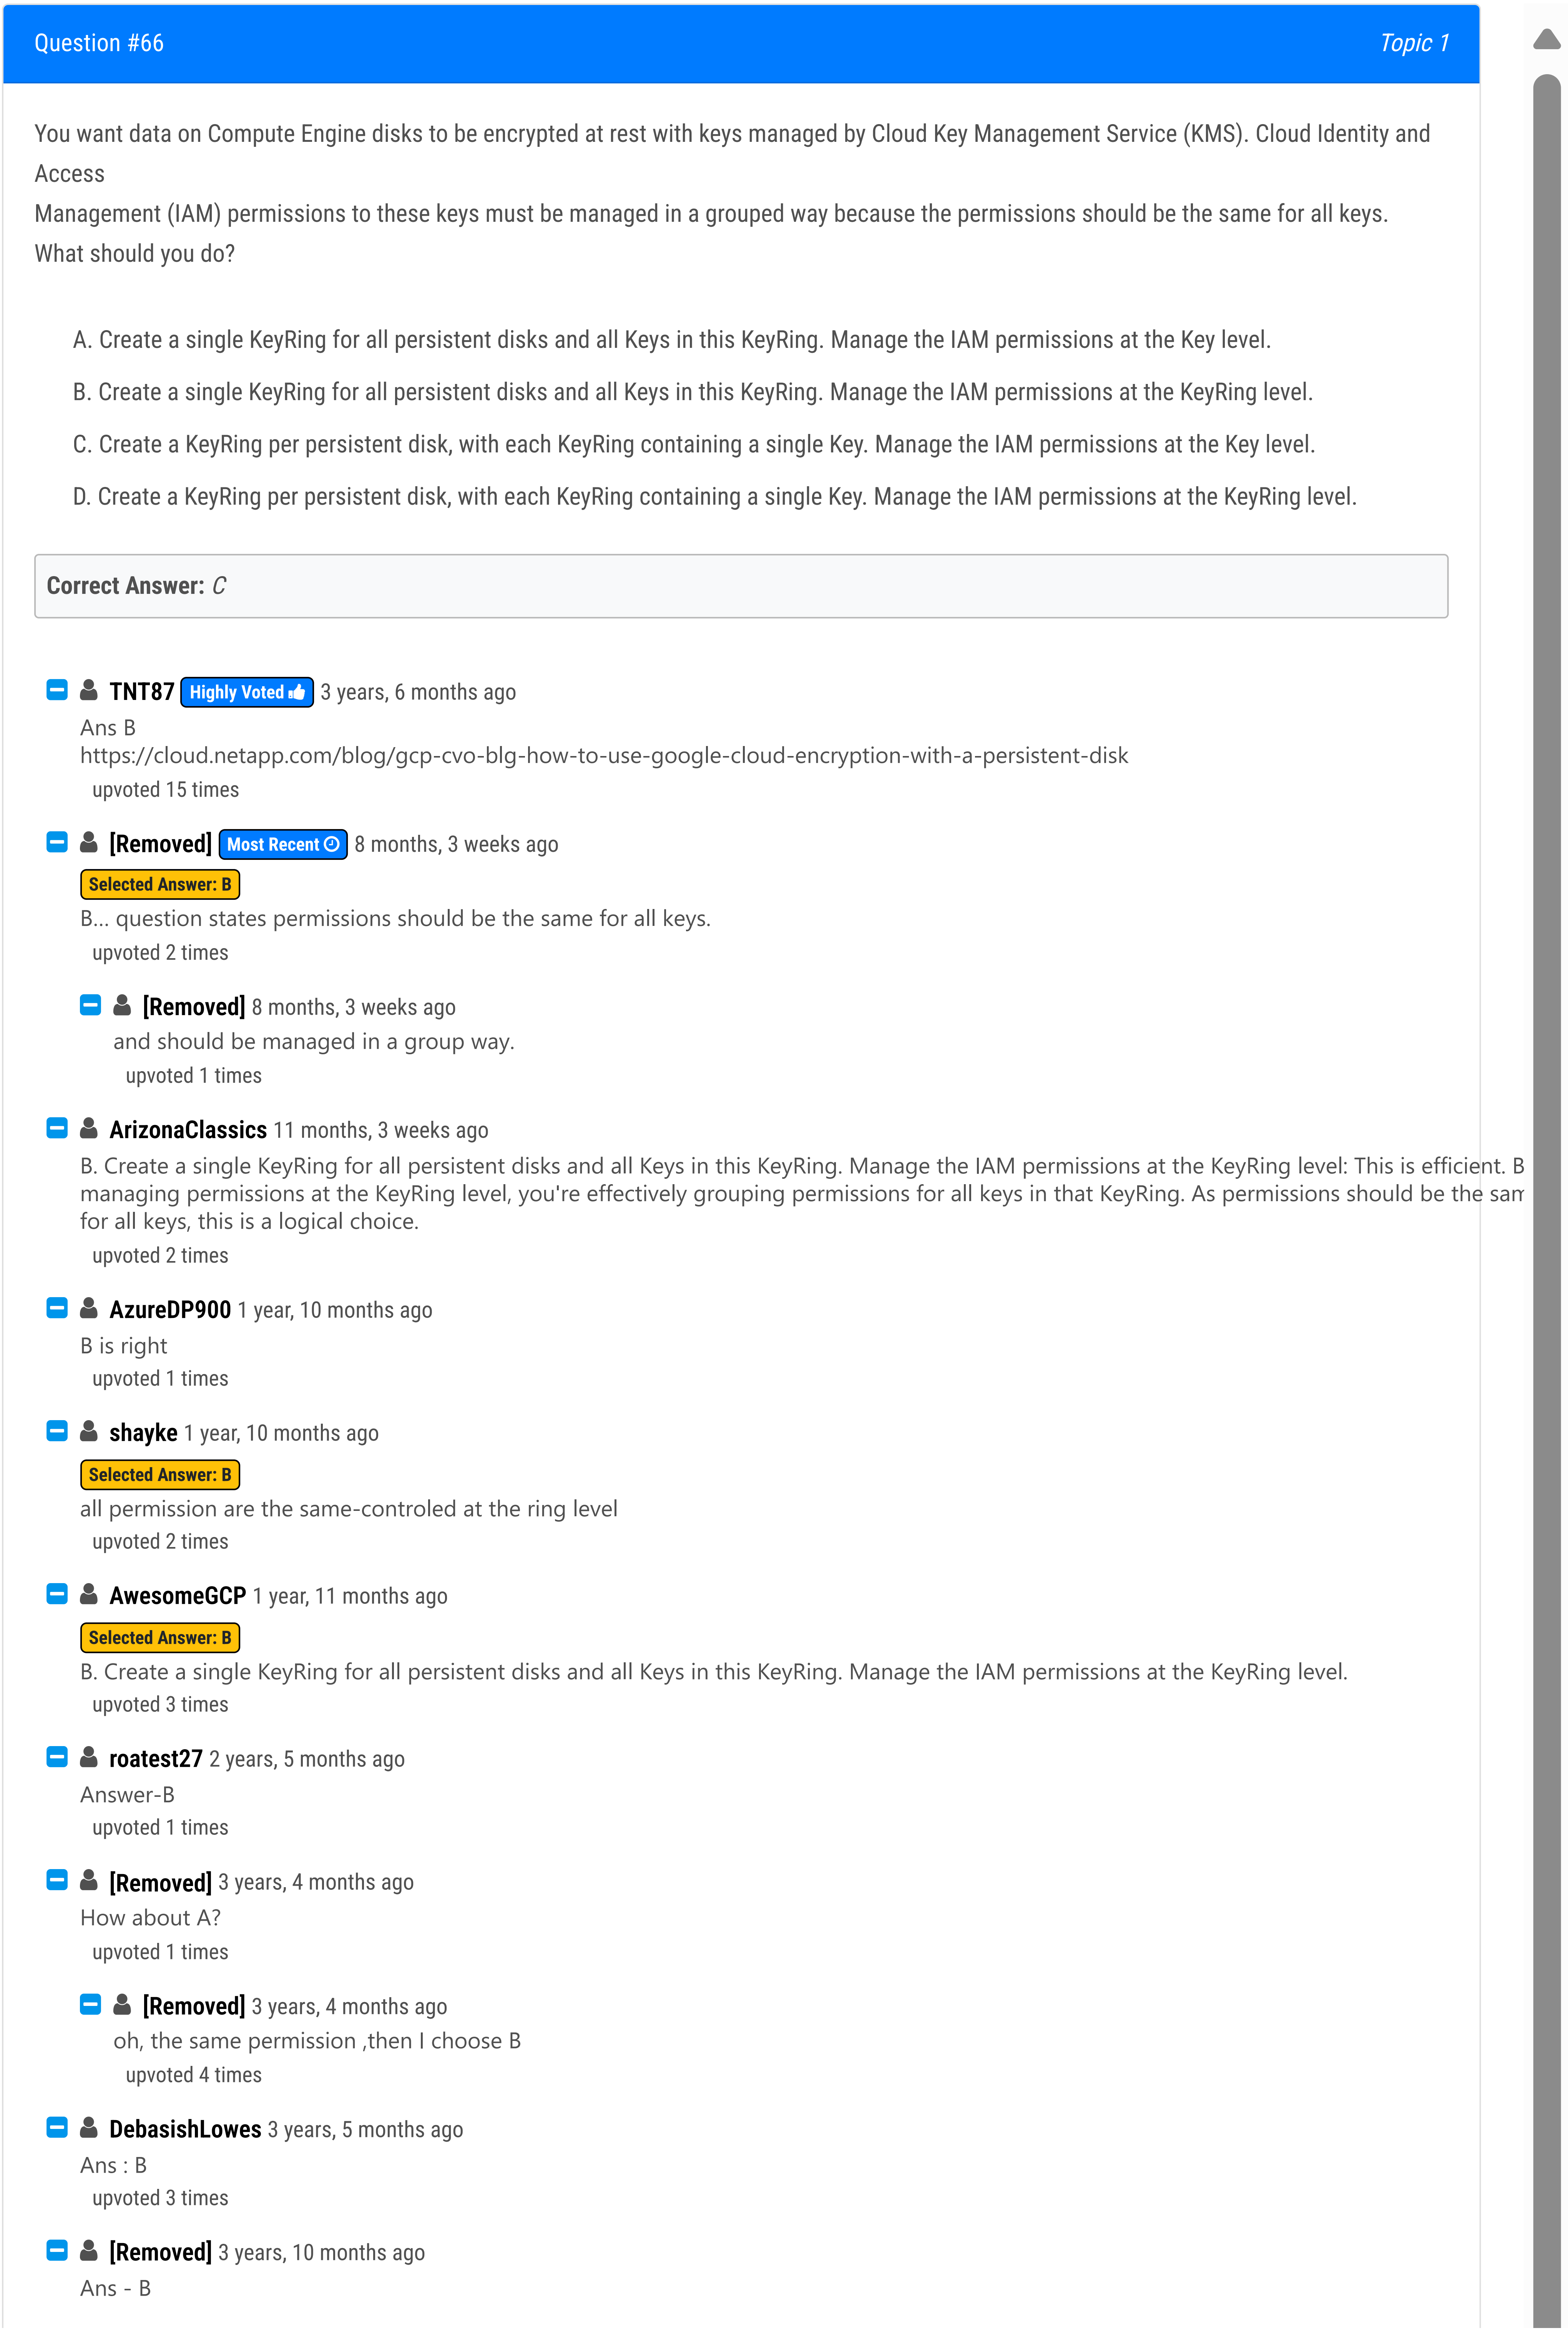

Question #66 *Topic 1*

You want data on Compute Engine disks to be encrypted at rest with keys managed by Cloud Key Management Service (KMS). Cloud Identity and Access
Management (IAM) permissions to these keys must be managed in a grouped way because the permissions should be the same for all keys.
What should you do?

A. Create a single KeyRing for all persistent disks and all Keys in this KeyRing. Manage the IAM permissions at the Key level.

B. Create a single KeyRing for all persistent disks and all Keys in this KeyRing. Manage the IAM permissions at the KeyRing level.

C. Create a KeyRing per persistent disk, with each KeyRing containing a single Key. Manage the IAM permissions at the Key level.

D. Create a KeyRing per persistent disk, with each KeyRing containing a single Key. Manage the IAM permissions at the KeyRing level.

Correct Answer: *C*

TNT87 Highly Voted 3 years, 6 months ago
Ans B
https://cloud.netapp.com/blog/gcp-cvo-blg-how-to-use-google-cloud-encryption-with-a-persistent-disk
upvoted 15 times

[Removed] Most Recent 8 months, 3 weeks ago
Selected Answer: B
B... question states permissions should be the same for all keys.
upvoted 2 times

[Removed] 8 months, 3 weeks ago
and should be managed in a group way.
upvoted 1 times

ArizonaClassics 11 months, 3 weeks ago
B. Create a single KeyRing for all persistent disks and all Keys in this KeyRing. Manage the IAM permissions at the KeyRing level: This is efficient. B managing permissions at the KeyRing level, you're effectively grouping permissions for all keys in that KeyRing. As permissions should be the sam for all keys, this is a logical choice.
upvoted 2 times

AzureDP900 1 year, 10 months ago
B is right
upvoted 1 times

shayke 1 year, 10 months ago
Selected Answer: B
all permission are the same-controled at the ring level
upvoted 2 times

AwesomeGCP 1 year, 11 months ago
Selected Answer: B
B. Create a single KeyRing for all persistent disks and all Keys in this KeyRing. Manage the IAM permissions at the KeyRing level.
upvoted 3 times

roatest27 2 years, 5 months ago
Answer-B
upvoted 1 times

[Removed] 3 years, 4 months ago
How about A?
upvoted 1 times

[Removed] 3 years, 4 months ago
oh, the same permission ,then I choose B
upvoted 4 times

DebasishLowes 3 years, 5 months ago
Ans : B
upvoted 3 times

[Removed] 3 years, 10 months ago
Ans - B

upvoted 1 times

Raushanr 3 years, 11 months ago

Answer-B

upvoted 1 times

Namaste 3 years, 11 months ago

B is the right answer

upvoted 1 times

MohitA 4 years ago

B should be the answer

upvoted 4 times

Question #67 *Topic 1*

A company is backing up application logs to a Cloud Storage bucket shared with both analysts and the administrator. Analysts should only have access to logs that do not contain any personally identifiable information (PII). Log files containing PII should be stored in another bucket that is only accessible by the administrator.
What should you do?

A. Use Cloud Pub/Sub and Cloud Functions to trigger a Data Loss Prevention scan every time a file is uploaded to the shared bucket. If the scan detects PII, have the function move into a Cloud Storage bucket only accessible by the administrator.

B. Upload the logs to both the shared bucket and the bucket only accessible by the administrator. Create a job trigger using the Cloud Data Loss Prevention API. Configure the trigger to delete any files from the shared bucket that contain PII.

C. On the bucket shared with both the analysts and the administrator, configure Object Lifecycle Management to delete objects that contain any PII.

D. On the bucket shared with both the analysts and the administrator, configure a Cloud Storage Trigger that is only triggered when PII data is uploaded. Use Cloud Functions to capture the trigger and delete such files.

Correct Answer: *C*

MohitA Highly Voted 4 years ago
A is the ans
upvoted 17 times

talktolanka Highly Voted 3 years, 4 months ago
Answer A
https://codelabs.developers.google.com/codelabs/cloud-storage-dlp-functions#0
https://www.youtube.com/watch?v=0TmO1f-Ox40
upvoted 8 times

Learn2fail Most Recent 11 months, 1 week ago
Selected Answer: A
A is answer
upvoted 2 times

AzureDP900 1 year, 10 months ago
A is right
upvoted 2 times

AwesomeGCP 1 year, 11 months ago
Selected Answer: A
A. Use Cloud Pub/Sub and Cloud Functions to trigger a Data Loss Prevention scan every time a file is uploaded to the shared bucket. If the scan detects PII, have the function move into a Cloud Storage bucket only accessible by theadministrator.
upvoted 4 times

[Removed] 1 year, 12 months ago
Selected Answer: A
A it is.
upvoted 2 times

[Removed] 3 years, 4 months ago
I also choose A.
upvoted 3 times

DebasishLowes 3 years, 5 months ago
Ans : A
upvoted 2 times

soukumar369 3 years, 8 months ago
Correct answer is A : Data Loss Prevention scan
upvoted 2 times

soukumar369 3 years, 8 months ago
A is correct.
upvoted 1 times

[Removed] 3 years, 10 months ago

Ans - A

upvoted 1 times

genesis3k 3 years, 10 months ago

Answer is A.

upvoted 1 times

passtest100 3 years, 11 months ago

SHOULD BE A

upvoted 1 times

Question #68 *Topic 1*

A customer terminates an engineer and needs to make sure the engineer's Google account is automatically deprovisioned.
What should the customer do?

A. Use the Cloud SDK with their directory service to remove their IAM permissions in Cloud Identity.

B. Use the Cloud SDK with their directory service to provision and deprovision users from Cloud Identity.

C. Configure Cloud Directory Sync with their directory service to provision and deprovision users from Cloud Identity.

D. Configure Cloud Directory Sync with their directory service to remove their IAM permissions in Cloud Identity.

Correct Answer: *C*

MohitA Highly Voted 4 years ago
C is the Answer
upvoted 7 times

ownez 3 years, 12 months ago
Agree with C.

"https://cloud.google.com/identity/solutions/automate-user-provisioning#cloud_identity_automated_provisioning"

"Cloud Identity has a catalog of automated provisioning connectors, which act as a bridge between Cloud Identity and third-party cloud apps.
upvoted 11 times

mynk29 2 years, 6 months ago
This option is for Cloud identity to third party app- you configure directory sync between AD and cloud identity.
upvoted 2 times

AzureDP900 1 year, 10 months ago
Agree with C, there is no need of cloud SDK.
upvoted 2 times

AzureDP900 1 year, 10 months ago
C. Configure Cloud Directory Sync with their directory service to provision and deprovision users from Cloud Identity.
upvoted 1 times

[Removed] Highly Voted 3 years, 10 months ago
Ans - C
upvoted 7 times

pradoUA Most Recent 11 months, 1 week ago
Selected Answer: C
C is correct
upvoted 2 times

AzureDP900 1 year, 10 months ago
C. Configure Cloud Directory Sync with their directory service to provision and deprovision users from Cloud Identity.
upvoted 1 times

AwesomeGCP 1 year, 11 months ago
Selected Answer: C
C. Configure Cloud Directory Sync with their directory service to provision and deprovision users from Cloud Identity.
upvoted 2 times

piyush_1982 2 years, 1 month ago
Selected Answer: C
Definitely C
upvoted 2 times

mynk29 2 years, 6 months ago
I don't think C is right answer. You configure Directory Sync to Sync from AD to cloud identity not the other way round.

Once a user is terminated- its account should be disabled on Directory and cloud identity will pick up via IAM. D looks more correct to me.
upvoted 2 times

AkbarM 1 year, 11 months ago

I also support D. The question may be provision and deprovision users. but technically it is to remove their IAM permissions in Cloud Identity. There is nothing like provision / deprovision user from cloud identity.

upvoted 1 times

rohan0411 8 months, 1 week ago

C is correct, because You cannot control IAM from Cloud Identity. Cloud identity only manages users and groups. It cannot remove IAM permissions through Cloud Identity.

upvoted 1 times

DebasishLowes 3 years, 5 months ago

Ans is C

upvoted 3 times

Question #69 *Topic 1*

An organization is evaluating the use of Google Cloud Platform (GCP) for certain IT workloads. A well-established directory service is used to manage user identities and lifecycle management. This directory service must continue for the organization to use as the `source of truth` directory for identities.
Which solution meets the organization's requirements?

A. Google Cloud Directory Sync (GCDS)

B. Cloud Identity

C. Security Assertion Markup Language (SAML)

D. Pub/Sub

Correct Answer: *B*
Reference:
https://cloud.google.com/solutions/federating-gcp-with-active-directory-introduction

desertlotus1211 Highly Voted 3 years, 5 months ago
The answer is A:
With Google Cloud Directory Sync (GCDS), you can synchronize the data in your Google Account with your Microsoft Active Directory or LDAP server. GCDS doesn't migrate any content (such as email messages, calendar events, or files) to your Google Account. You use GCDS to synchroni your Google users, groups, and shared contacts to match the information in your LDAP server.

The questions says the well established directory service is the 'source of truth' not GCP... So LDAP or AD is the source... GCDS will sync that to match those, not replace them...
upvoted 17 times

AzureDP900 1 year, 10 months ago
Agreed
upvoted 2 times

subhala Highly Voted 3 years, 9 months ago
GCDS -? It helps sync up from the source of truth (any IdP like ldap, AD) to Google identity. In this scenario, the question is what can be a good identity service by itself, hence B is the right answer.
upvoted 12 times

desertlotus1211 1 year ago
The question inplies the company has a directory as the soruce of truth and want to maintain that in GCP... GCDS will make sure that occurs tc Cloud Identity. It's not askling for a replacement of LDAP/AD.
upvoted 2 times

ArizonaClassics Most Recent 11 months, 3 weeks ago
Google Cloud Directory Sync (GCDS): GCDS is a tool used to synchronize your Google Workspace user data with your Microsoft Active Directory c other LDAP servers. This would ensure that Google Workspace has the same user data as your existing directory, but it doesn't act as an identity provider (IDP).
BUT

C. Security Assertion Markup Language (SAML): SAML is an open standard for exchanging authentication and authorization data between an identity provider (your organization's existing directory service) and a service provider (like GCP). With SAML, GCP can rely on your existing directory service for authentication, and your existing directory remains the "source of truth."
upvoted 2 times

PST21 1 year, 8 months ago
Orgn is evaluating GC so cloud Identity is the GC product hence B
upvoted 1 times

AwesomeGCP 1 year, 11 months ago
Selected Answer: A
A. Google Cloud Directory Sync (GCDS)
upvoted 4 times

cloudprincipal 2 years, 3 months ago
Selected Answer: A
With Google Cloud Directory Sync (GCDS), you can synchronize the data in your Google Account with your Microsoft Active Directory or LDAP server. GCDS doesn't migrate any content (such as email messages, calendar events, or files) to your Google Account. You use GCDS to synchroni your Google users, groups, and shared contacts to match the information in your LDAP server.
https://support.google.com/a/answer/106368?hl=en
upvoted 3 times

szl0144 2 years, 3 months ago

B should be the answer, GCDS is for ad sync.

upvoted 2 times

MariaGabiGabriela 2 years, 3 months ago

Yes, but identity by itself will solve nothing, the user would have to recreate all users and thus have a different IDP, this clearly goes against the question

upvoted 2 times

Bill831231 2 years, 8 months ago

seems there is nothing metioned about what they have on premise, so B is better

upvoted 1 times

syllox 3 years, 4 months ago

Answer A

upvoted 3 times

WakandaF 3 years, 4 months ago

A or B?

upvoted 2 times

DebasishLowes 3 years, 5 months ago

Ans : B as per the question.

upvoted 1 times

asee 3 years, 6 months ago

My Answer will go for A (GCDS), noticed the question is mentioning about "A directory service 'is used' " / "must continue" instead of "A directory service 'will be used' ". So here my understanding is the organization has already using their own directory service. Therefore Answer B - Cloud identity may not be an option.

upvoted 4 times

KWatHK 3 years, 7 months ago

Ans is B because the questions said "the well-established directory must continue for the orgnanization to use as the source of truth" so that the user access to GCP must authenticated by the existing directory. Cloud Identity support to federate it to 3rd party/ADFS using SAML.

upvoted 1 times

mikelabs 3 years, 9 months ago

GCDS is an app to sync users, groups and other features from AD to Cloud Identity. But, in this question, the customer needs to know what's the product on GCP that meet with this. So, I thiink the answer is B.

upvoted 8 times

[Removed] 3 years, 10 months ago

Ans - A

upvoted 3 times

ownez 3 years, 12 months ago

GCDS is a part of CI's feature that synchronizes the data in Google domain to match with AD/LDAP server. This includes users, groups contacts et are synchronized/migrated to match.

Hence, I would go B.

"https://se-cloud-experts.com/wp/wp-content/themes/se-it/images/pdf/google-cloud-identity-services.pdf"

upvoted 3 times

ownez 3 years, 11 months ago

Sorry. It's A.

upvoted 2 times

bogdant 4 years ago

Isn't it A?

upvoted 2 times

Sheeda 4 years ago

That is used to sync, not the directly itself

upvoted 1 times

Fellipo 3 years, 9 months ago

A well-established directory service , so "A"

upvoted 2 times

MohitA 4 years ago

Agree A

upvoted 4 times

Question #70 *Topic 1*

Which international compliance standard provides guidelines for information security controls applicable to the provision and use of cloud services?

A. ISO 27001

B. ISO 27002

C. ISO 27017

D. ISO 27018

Correct Answer: *C*
Create a new Service Account that should be able to list the Compute Engine instances in the project. You want to follow Google-recommended practices.

asee Highly Voted 👍 3 years, 6 months ago
Yes, My answer also goes to C and my last compliance related project is also working on ISO27017 in order to extend the scope to Cloud service user/provider.
upvoted 11 times

AzureDP900 1 year, 10 months ago
C is right
upvoted 1 times

AzureDP900 1 year, 10 months ago
https://cloud.google.com/security/compliance/iso-27017
upvoted 2 times

pradoUA Most Recent ⏱ 11 months, 1 week ago
Selected Answer: C
C. ISO 27017
upvoted 2 times

AwesomeGCP 1 year, 11 months ago
Selected Answer: C
C. ISO 27017
upvoted 4 times

certificationjjmmm 2 years, 1 month ago
C is correct.
https://cloud.google.com/security/compliance/iso-27017
upvoted 3 times

[Removed] 3 years, 10 months ago
Ans - C
upvoted 3 times

Namaste 3 years, 11 months ago
CCSP Question...C is the Answer
upvoted 3 times

ownez 4 years ago
C is correct.

"https://www.iso.org/standard/43757.html"
upvoted 4 times

Question #71 *Topic 1*

You will create a new Service Account that should be able to list the Compute Engine instances in the project. You want to follow Google-recommended practices.
What should you do?

A. Create an Instance Template, and allow the Service Account Read Only access for the Compute Engine Access Scope.

B. Create a custom role with the permission compute.instances.list and grant the Service Account this role.

C. Give the Service Account the role of Compute Viewer, and use the new Service Account for all instances.

D. Give the Service Account the role of Project Viewer, and use the new Service Account for all instances.

Correct Answer: *A*

MohitA Highly Voted 4 years ago
B, https://cloud.google.com/compute/docs/access/iam
upvoted 16 times

mlyu 4 years ago
Although it is not encourage to use custome role, but last sentence in the answer C makes B be the only option
upvoted 6 times

AzureDP900 1 year, 10 months ago
B is right
upvoted 2 times

Roflcopter Highly Voted 2 years ago
Selected Answer: B
Key here is "and grant the Service Account this role.". C and D are giving this role to ALL instances which is overly permissive. A is wrong. Only choice is B
upvoted 5 times

ArizonaClassics Most Recent 11 months, 3 weeks ago
B. Create a custom role with the permission compute.instances.list and grant the Service Account this role: This follows the principle of least privilege by granting only the specific permission needed.
upvoted 2 times

Brosh 1 year, 8 months ago
I don't get why is it not C, you grant that specific service account the role over all instances, is it wrong because that service account will be able t view not only compute instances?
upvoted 2 times

shayke 1 year, 8 months ago
Selected Answer: B
B is the right ans - you only want to list the instances
upvoted 3 times

Meyucho 1 year, 8 months ago
Selected Answer: B
With C the SA will list ONLY the instances that are configured to use that SA.
The option B will give permissions to list ALL instances.
upvoted 3 times

AwesomeGCP 1 year, 11 months ago
Selected Answer: B
B. Create a custom role with the permission compute.instances.list and grant the Service Account this role.
upvoted 3 times

nbrnschwgr 2 years ago
C. because google recommends pre-defined narrow scope roles over custom roles.
upvoted 2 times

cloudprincipal 2 years, 3 months ago
Selected Answer: B
The roles/compute.viewer provides a lot more privileges than just listing compute instances
upvoted 4 times

cloudprincipal 2 years, 3 months ago

Selected Answer: C

Compute Viewer
Read-only access to get and list Compute Engine resources, without being able to read the data stored on them.

https://cloud.google.com/compute/docs/access/iam#compute.viewer

upvoted 2 times

> **cloudprincipal** 2 years, 2 months ago
>
> This is incorrect, as Compute Viewer provides a lot more than what is required
>
> upvoted 1 times

sudarchary 2 years, 7 months ago

B. The only option that adheres to the principle of least privilege and meets
question requirements is B

upvoted 5 times

[Removed] 3 years, 4 months ago

I think C is good

upvoted 4 times

DebasishLowes 3 years, 5 months ago

Ans : B

upvoted 1 times

dtmtor 3 years, 5 months ago

Ans is B

upvoted 1 times

[Removed] 3 years, 10 months ago

Ans - B

upvoted 1 times

genesis3k 3 years, 10 months ago

Answer is B, based on least privilege principle.

upvoted 1 times

Question #72 *Topic 1*

In a shared security responsibility model for IaaS, which two layers of the stack does the customer share responsibility for? (Choose two.)

A. Hardware

B. Network Security

C. Storage Encryption

D. Access Policies

E. Boot

Correct Answer: *CD*

DebasishLowes Highly Voted 3 years, 5 months ago

Ans : BD

upvoted 12 times

AliHammoud Most Recent 5 months, 2 weeks ago

B and D

upvoted 1 times

GCBC 1 year ago

look at diagram, its B D -> https://cloud.google.com/architecture/framework/security/shared-responsibility-shared-fate#shared-diagram

upvoted 4 times

GCBC 1 year ago

B. Network Security
D. Access Policies

upvoted 2 times

sushmitha95 1 year, 7 months ago

Selected Answer: BD

D. Access Policies B. Network Security

upvoted 3 times

shayke 1 year, 8 months ago

b and D - according to the shared responsibility moder for IAAS

upvoted 2 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: BD

B. Network Security
D. Access Policies

upvoted 3 times

Random_Mane 1 year, 11 months ago

Selected Answer: BD

Chart is here https://cloud.google.com/architecture/framework/security/shared-responsibility-shared-fate

upvoted 3 times

rr4444 2 years, 8 months ago

Selected Answer: BD

BD https://cloud.google.com/blog/products/containers-kubernetes/exploring-container-security-the-shared-responsibility-model-in-gke-container-security-shared-responsibility-model-gke

upvoted 4 times

[Removed] 3 years, 10 months ago

Ans - BD

upvoted 4 times

saurabh1805 3 years, 10 months ago

B and D is correct option.

upvoted 4 times

passtest100 3 years, 11 months ago

B and D

upvoted 4 times

lordb 3 years, 11 months ago

B and D

upvoted 4 times

Question #73 *Topic 1*

An organization is starting to move its infrastructure from its on-premises environment to Google Cloud Platform (GCP). The first step the organization wants to take is to migrate its ongoing data backup and disaster recovery solutions to GCP. The organization's on-premises production environment is going to be the next phase for migration to GCP. Stable networking connectivity between the on-premises environment and GCP is also being implemented.
Which GCP solution should the organization use?

A. BigQuery using a data pipeline job with continuous updates via Cloud VPN

B. Cloud Storage using a scheduled task and gsutil via Cloud Interconnect

C. Compute Engines Virtual Machines using Persistent Disk via Cloud Interconnect

D. Cloud Datastore using regularly scheduled batch upload jobs via Cloud VPN

Correct Answer: *B*
Reference:
https://cloud.google.com/solutions/migration-to-google-cloud-building-your-foundation

ownez Highly Voted 3 years, 11 months ago
Agree B.

https://cloud.google.com/solutions/dr-scenarios-for-data#production_environment_is_on-premises
upvoted 11 times

madcloud32 Most Recent 6 months ago
Selected Answer: B
Data Backup to GCP, so B is correct
upvoted 1 times

Xoxoo 11 months, 3 weeks ago
Selected Answer: B
To migrate ongoing data backup and disaster recovery solutions to Google Cloud Platform (GCP), the most suitable GCP solution for the organization would be Cloud Storage using a scheduled task and gsutil via Cloud Interconnect. This solution offers scalability, cost-efficiency, and features essential for backup and disaster recovery solutions.

Cloud Storage provides a scalable object storage service that allows you to store and retrieve large amounts of data. By using a scheduled task ar gsutil, you can automate the backup process and ensure that your data is securely stored in the cloud. Cloud Interconnect ensures stable networking connectivity between the on-premises environment and GCP, making it an ideal choice for migrating data backup and disaster recovery solutions
upvoted 3 times

TNT87 1 year, 5 months ago
https://cloud.google.com/architecture/dr-scenarios-for-data#back-up-to-cloud-storage-using-a-scheduled-task
upvoted 1 times

shayke 1 year, 8 months ago
Selected Answer: B
B- backup and DR is GCS
upvoted 2 times

rotorclear 1 year, 10 months ago
Selected Answer: B
https://medium.com/@pvergadia/cold-disaster-recovery-on-google-cloud-for-applications-running-on-premises-114b31933d02
upvoted 2 times

AzureDP900 1 year, 10 months ago
B is correct
upvoted 1 times

AwesomeGCP 1 year, 11 months ago
Selected Answer: B
B. Cloud Storage using a scheduled task and gsutil via Cloud Interconnect
upvoted 2 times

cloudprincipal 2 years, 3 months ago
Selected Answer: B

https://cloud.google.com/solutions/dr-scenarios-for-data#production_environment_is_on-premises

upvoted 1 times

rr4444 2 years, 8 months ago

Selected Answer: C

Disaster recover made me think C Compute Engines Virtual Machines using Persistent Disk via Cloud Interconnect

Disaster recovery with remote backup alone, when all prod is on premise, will take too long to be viable. The VMs don't need to be running when no disaster

upvoted 3 times

csrazdan 1 year, 9 months ago

You would have been correct if the question had any RTO/RPO specifications. In absence of this question is assuming backup and restore as a DR strategy. So Option B Cloud Storage is the correct answer.

upvoted 1 times

desertlotus1211 1 year ago

You never move compute first...

upvoted 1 times

DebasishLowes 3 years, 5 months ago

Ans : B

upvoted 2 times

[Removed] 3 years, 10 months ago

Ans - V

upvoted 1 times

[Removed] 3 years, 10 months ago

Typo - it's B

upvoted 2 times

Question #74 *Topic 1*

What are the steps to encrypt data using envelope encryption?

A.

✑ Generate a data encryption key (DEK) locally.

✑ Use a key encryption key (KEK) to wrap the DEK.

✑ Encrypt data with the KEK.

✑ Store the encrypted data and the wrapped KEK.

B.

✑ Generate a key encryption key (KEK) locally.

✑ Use the KEK to generate a data encryption key (DEK).

✑ Encrypt data with the DEK.

✑ Store the encrypted data and the wrapped DEK.

C.

✑ Generate a data encryption key (DEK) locally.

✑ Encrypt data with the DEK.

✑ Use a key encryption key (KEK) to wrap the DEK.

✑ Store the encrypted data and the wrapped DEK.

D.

✑ Generate a key encryption key (KEK) locally.

✑ Generate a data encryption key (DEK) locally.

✑ Encrypt data with the KEK.

Store the encrypted data and the wrapped DEK.

▪

Correct Answer: *C*

Reference:

https://cloud.google.com/kms/docs/envelope-encryption

Tabayashi Highly Voted 2 years, 4 months ago

Answer is (C).

The process of encrypting data is to generate a DEK locally, encrypt data with the DEK, use a KEK to wrap the DEK, and then store the encrypted data and the wrapped DEK. The KEK never leaves Cloud KMS.
https://cloud.google.com/kms/docs/envelope-encryption#how_to_encrypt_data_using_envelope_encryption

upvoted 19 times

AzureDP900 1 year, 10 months ago

C is right

upvoted 3 times

Mr_MIXER007 Most Recent 4 days, 17 hours ago

Answer is (C).

upvoted 1 times

desertlotus1211 1 year ago

Answer is C;

https://cloud.google.com/kms/docs/envelope-encryption#:~:text=decrypt%20data%20directly.-,How%20to%20encrypt%20data%20using%20envelope%20encryption,data%20and%20the%20v
apped%20DEK.

upvoted 3 times

Appsec977 1 year, 3 months ago

C is the correct solution because KEK is never generated on the client's side, KEK is stored in GCP.

upvoted 4 times

AwesomeGCP 1 year, 11 months ago

Answer - C is correct.
https://cloud.google.com/kms/docs/envelope-encryption#how_to_encrypt_data_using_envelope_encryption

upvoted 3 times

[Removed] 1 year, 12 months ago

C it is

upvoted 3 times

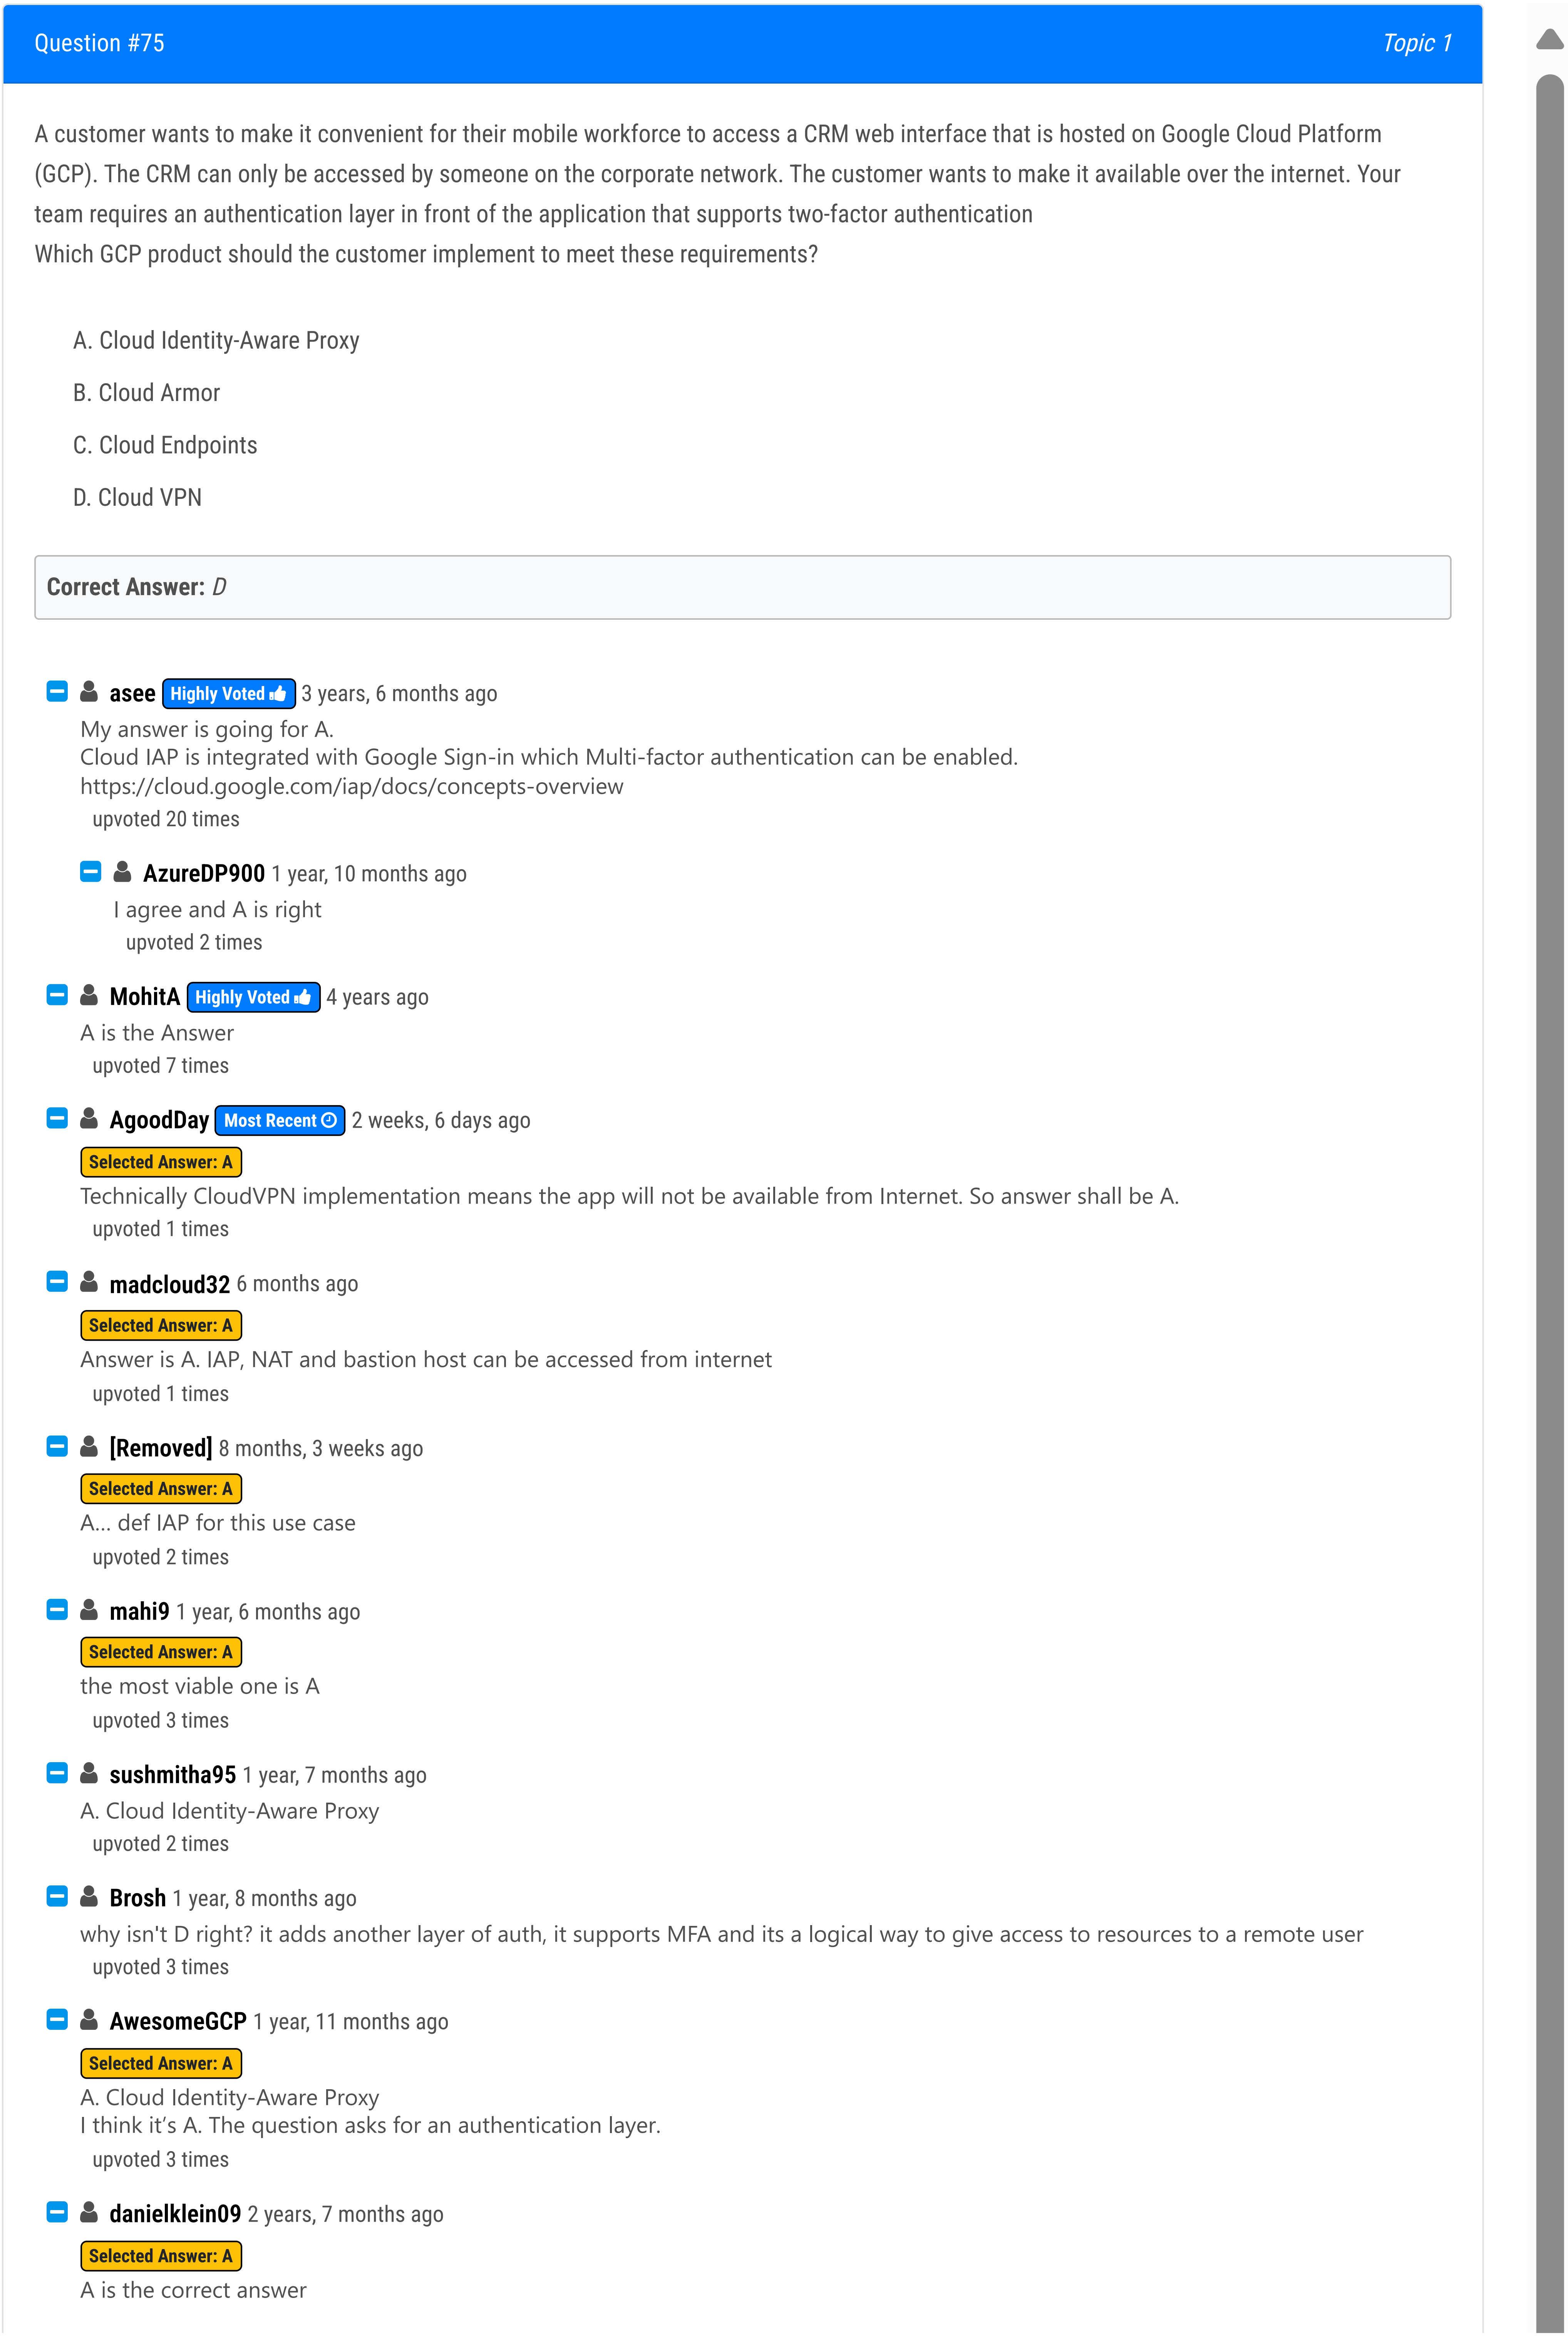

Question #75 *Topic 1*

A customer wants to make it convenient for their mobile workforce to access a CRM web interface that is hosted on Google Cloud Platform (GCP). The CRM can only be accessed by someone on the corporate network. The customer wants to make it available over the internet. Your team requires an authentication layer in front of the application that supports two-factor authentication
Which GCP product should the customer implement to meet these requirements?

A. Cloud Identity-Aware Proxy

B. Cloud Armor

C. Cloud Endpoints

D. Cloud VPN

Correct Answer: *D*

asee Highly Voted 3 years, 6 months ago
My answer is going for A.
Cloud IAP is integrated with Google Sign-in which Multi-factor authentication can be enabled.
https://cloud.google.com/iap/docs/concepts-overview
upvoted 20 times

AzureDP900 1 year, 10 months ago
I agree and A is right
upvoted 2 times

MohitA Highly Voted 4 years ago
A is the Answer
upvoted 7 times

AgoodDay Most Recent 2 weeks, 6 days ago
Selected Answer: A
Technically CloudVPN implementation means the app will not be available from Internet. So answer shall be A.
upvoted 1 times

madcloud32 6 months ago
Selected Answer: A
Answer is A. IAP, NAT and bastion host can be accessed from internet
upvoted 1 times

[Removed] 8 months, 3 weeks ago
Selected Answer: A
A... def IAP for this use case
upvoted 2 times

mahi9 1 year, 6 months ago
Selected Answer: A
the most viable one is A
upvoted 3 times

sushmitha95 1 year, 7 months ago
A. Cloud Identity-Aware Proxy
upvoted 2 times

Brosh 1 year, 8 months ago
why isn't D right? it adds another layer of auth, it supports MFA and its a logical way to give access to resources to a remote user
upvoted 3 times

AwesomeGCP 1 year, 11 months ago
Selected Answer: A
A. Cloud Identity-Aware Proxy
I think it’s A. The question asks for an authentication layer.
upvoted 3 times

danielklein09 2 years, 7 months ago
Selected Answer: A
A is the correct answer

upvoted 3 times

[Removed] 3 years, 10 months ago

Ans - A

upvoted 4 times

passtest100 3 years, 11 months ago

SHOULD BE A

upvoted 5 times

Raushanr 3 years, 11 months ago

Answer -A

upvoted 4 times

Question #76 *Topic 1*

Your company is storing sensitive data in Cloud Storage. You want a key generated on-premises to be used in the encryption process.
What should you do?

A. Use the Cloud Key Management Service to manage a data encryption key (DEK).

B. Use the Cloud Key Management Service to manage a key encryption key (KEK).

C. Use customer-supplied encryption keys to manage the data encryption key (DEK).

D. Use customer-supplied encryption keys to manage the key encryption key (KEK).

Correct Answer: *A*
Reference:
https://cloud.google.com/security/encryption-at-rest/default-encryption/

HateMicrosoft Highly Voted 👍 3 years, 5 months ago

The anwser is:C
This is a Customer-supplied encryption keys (CSEK).
We generate our own encryption key and manage it on-premises.
A KEK never leaves Cloud KMS.There is no KEK or KMS on-premises.

Encryption at rest by default, with various key management options
https://cloud.google.com/security/encryption-at-rest

upvoted 31 times

sudarchary Highly Voted 👍 2 years, 7 months ago

Selected Answer: D

Reference Links:
https://cloud.google.com/kms/docs/envelope-encryption
https://cloud.google.com/security/encryption-at-rest/customer-supplied-encryption-keys

upvoted 9 times

Mr_MIXER007 Most Recent ⏱ 4 days, 17 hours ago

Selected Answer: C

The anwser is:C

upvoted 1 times

3d9563b 1 month, 1 week ago

Selected Answer: C

By using customer-supplied encryption keys (CSEK) to manage the data encryption key (DEK), you can ensure that the encryption process utilizes key that was generated and controlled on-premises, meeting your security and compliance requirements.

upvoted 1 times

salamKvelas 3 months, 2 weeks ago

`customer-supplied encryption keys` == `DEK`, so the only answer that makes sense is A use KMS for KEK to wrap the DEK

upvoted 1 times

shanwford 4 months ago

Selected Answer: C

Can't be A/B because "key generated on-premises" requirement. KEK ist KMS specific.
Why (C):
https://cloud.google.com/docs/security/encryption/customer-supplied-encryption-keys#cloud_storage --> "The raw CSEK is used to unwrap wrapped chunk keys, to create raw chunk keys in memory. These are used to decrypt data chunks stored in the storage systems. These keys are used as the data encryption keys (DEK) in Google Cloud Storage for your data."

upvoted 1 times

madcloud32 6 months ago

Selected Answer: C

C is answer. DEK

upvoted 1 times

mjcts 6 months, 4 weeks ago

Selected Answer: C

Customer-supplied because it is generated on prem. And we can only talk about DEK. KEK is always managed by Google

upvoted 1 times

rsamant 9 months ago

D , CSEK is used for KEK , DEK is always generated by Google as different chunks use different DEK

Raw CSEK Storage system memory Provided by the customer.
Key encryption key (KEK) for chunk keys.
Wraps the chunk keys. Customer-requested operation (e.g., insertObject or getObject) is complete

https://cloud.google.com/docs/security/encryption/customer-supplied-encryption-keys

upvoted 3 times

rottzy 11 months, 2 weeks ago

C, KEK is google managed

upvoted 1 times

Xoxoo 11 months, 2 weeks ago

Selected Answer: C

To use a key generated on-premises for encrypting data in Cloud Storage, you should:

C. Use customer-supplied encryption keys to manage the data encryption key (DEK).

With customer-supplied encryption keys (CSEK), you can provide your own encryption keys, generated and managed on-premises, to encrypt and decrypt data in Cloud Storage. The data encryption key (DEK) is the key used to encrypt the actual data, and by using CSEK, you can manage this key with your own on-premises key management system.

upvoted 1 times

Xoxoo 11 months, 2 weeks ago

Options A and B involve using Google Cloud's Key Management Service (KMS), which generates and manages encryption keys within Google Cloud, not on-premises.

Option D is not a common practice and is not directly supported for encrypting data in Cloud Storage.

upvoted 2 times

ananta93 12 months ago

Selected Answer: C

The Answer is C. The raw CSEK is used to unwrap wrapped chunk keys, to create raw chunk keys in memory. These are used to decrypt data chunk stored in the storage systems. These keys are used as the data encryption keys (DEK) in Google Cloud Storage for your data.

https://cloud.google.com/docs/security/encryption/customer-supplied-encryption-keys#cloud_storage

upvoted 2 times

desertlotus1211 1 year ago

Answer is C:
https://cloud.google.com/docs/security/encryption/customer-supplied-encryption-keys#cloud_storage

If you look at the ENTIRE process - it CSEK is used to create the DEK (final product) for decryption if its data...

upvoted 3 times

RuchiMishra 1 year ago

Selected Answer: D

https://cloud.google.com/docs/security/encryption/customer-supplied-encryption-keys#cloud_storage

upvoted 2 times

civilizador 1 year, 1 month ago

C . The answer is C and I don't understand why some people here rewriting google official doc here and saying answer is D?? Here is the link please read it carefully this is not an Instagramm feed. Please when you reading 3 seconds and come here you start confusing many people . Here is link SPECIFICALLY FOR CLOUD STORAGE . https://cloud.google.com/docs/security/encryption/customer-supplied-encryption-keys#cloud_storage

upvoted 3 times

MaryKey 1 year ago

I'm confused here - the article on Google says literally:
"Raw CSEK - Provided by the customer.
Key encryption key (KEK) for chunk keys.
Wraps the chunk keys".
In other words - KEK, not DEK

upvoted 3 times

[Removed] 1 year, 1 month ago

Selected Answer: C

"C"
KEK never leaves Cloud KMS.
Customer supplied key can only be for DEK.

upvoted 3 times

Ishu_awsguy 1 year, 2 months ago

D is the correct answer .
When we do CSEk that onprem key acts as the KEK.

upvoted 1 times

Question #77 *Topic 1*

Last week, a company deployed a new App Engine application that writes logs to BigQuery. No other workloads are running in the project. You need to validate that all data written to BigQuery was done using the App Engine Default Service Account.
What should you do?

A. 1. Use Cloud Logging and filter on BigQuery Insert Jobs. 2. Click on the email address in line with the App Engine Default Service Account in the authentication field. 3. Click Hide Matching Entries. 4. Make sure the resulting list is empty.

B. 1. Use Cloud Logging and filter on BigQuery Insert Jobs. 2. Click on the email address in line with the App Engine Default Service Account in the authentication field. 3. Click Show Matching Entries. 4. Make sure the resulting list is empty.

C. 1. In BigQuery, select the related dataset. 2. Make sure that the App Engine Default Service Account is the only account that can write to the dataset.

D. 1. Go to the Identity and Access Management (IAM) section of the project. 2. Validate that the App Engine Default Service Account is the only account that has a role that can write to BigQuery.

Correct Answer: *C*

AwesomeGCP Highly Voted 1 year, 11 months ago

Selected Answer: A

A. 1. Use StackDriver Logging and filter on BigQuery Insert Jobs.
2. Click on the email address in line with the App Engine Default Service Account in the authentication field.
3. Click Hide Matching Entries.
4. Make sure the resulting list is empty.

upvoted 13 times

> **Appsec977** 1 year, 3 months ago
>
> Stackdriver is now Cloud Operations.
>
> upvoted 2 times

blacortik Highly Voted 1 year ago

Selected Answer: B

A: This option seems to be about using Cloud Logging and hiding matching entries. However, hiding matching entries wouldn't help in verifying the specific service account used for BigQuery Insert Jobs.
C: While restricting permissions in BigQuery is important for security, it doesn't directly help you validate the specific service account that wrote tl data.
D: While IAM roles and permissions are important to manage access, it doesn't provide a clear process for verifying the service account used for a specific action.

In summary, option B provides the appropriate steps to validate that data written to BigQuery was done using the App Engine Default Service Account by examining the Cloud Logging entries.

upvoted 5 times

dija123 Most Recent 5 months, 2 weeks ago

Selected Answer: B

Agree with B

upvoted 1 times

> **dija123** 5 months, 1 week ago
>
> I think "Make sure the resulting list is empty" makes answer A is correct not B
>
> upvoted 3 times

PST21 1 year, 8 months ago

A is correct as last 2 are means of doing it rather than validating it

upvoted 2 times

shayke 1 year, 10 months ago

Selected Answer: C

validate - C

upvoted 1 times

tangac 1 year, 12 months ago

Selected Answer: A

https://www.examtopics.com/discussions/google/view/32259-exam-professional-cloud-security-engineer-topic-1-question/

upvoted 4 times

Question #78 *Topic 1*

Your team wants to limit users with administrative privileges at the organization level.
Which two roles should your team restrict? (Choose two.)

A. Organization Administrator

B. Super Admin

C. GKE Cluster Admin

D. Compute Admin

E. Organization Role Viewer

Correct Answer: *AB*
Reference:
https://cloud.google.com/resource-manager/docs/creating-managing-organization

HateMicrosoft Highly Voted 3 years, 5 months ago
The correct anwser is : A&B
-resourcemanager.organizationAdmin
-Cloud Identity super admin(Old G-Suite Google Workspace)
upvoted 14 times

[Removed] Most Recent 8 months, 3 weeks ago
Selected Answer: AD
For me the correct answer A & D. In the context of gcp there is no super admin. Super admin is only used in gsuite.
upvoted 2 times

AwesomeGCP 1 year, 11 months ago
Selected Answer: AB
A. Organization Administrator
B. Super Admin
upvoted 4 times

> **AzureDP900** 1 year, 10 months ago
> AB is correct
> upvoted 1 times

Bingo21 3 years, 6 months ago
It says "limit users with administrative privileges" - D doesnt give you admin privileges. AB is the closest to what the question is looking for.
upvoted 3 times

[Removed] 3 years, 10 months ago
Ans - AB
upvoted 3 times

MohitA 4 years ago
AB are the one
upvoted 4 times

> **singhjoga** 3 years, 8 months ago
> There is no such role as "Super Admin". There is a Super Admin user. which has the "Owner" role to the how Organisation.
> Answer is probably A and D.
> upvoted 8 times

Question #79 *Topic 1*

An organization's security and risk management teams are concerned about where their responsibility lies for certain production workloads they are running in
Google Cloud and where Google's responsibility lies. They are mostly running workloads using Google Cloud's platform-as-a-Service (PaaS) offerings, including
App Engine primarily.
Which area in the technology stack should they focus on as their primary responsibility when using App Engine?

A. Configuring and monitoring VPC Flow Logs

B. Defending against XSS and SQLi attacks

C. Managing the latest updates and security patches for the Guest OS

D. Encrypting all stored data

Correct Answer: *D*

Random_Mane Highly Voted 1 year, 11 months ago

Selected Answer: B

B. in PaaS the customer is responsible for web app security, deployment, usage, access policy, and content.
https://cloud.google.com/architecture/framework/security/shared-responsibility-shared-fate

upvoted 7 times

madcloud32 Most Recent 6 months ago

Selected Answer: B

B is correct. Defense of App Engine and Application Security.

upvoted 1 times

gcpengineer 1 year, 3 months ago

Selected Answer: B

B is the ans.

upvoted 2 times

AzureDP900 1 year, 10 months ago

B is correct

upvoted 2 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: B

B. Defending against XSS and SQLi attacks
Data at rest is encrypted by default by Google. So D is wrong. Should be B.

upvoted 4 times

koko2314 1 year, 11 months ago

Answer should be D. For SAAS solutions web based attacks are managed by Google. We just need to take care of the data as per the link below.

upvoted 1 times

> **desertlotus1211** 1 year ago
>
> read the question again... it's not D
>
> upvoted 1 times

GHOST1985 1 year, 11 months ago

Selected Answer: D

Answer is D

upvoted 1 times

> **GHOST1985** 1 year, 11 months ago
>
> In PaaS, we're responsible for more controls than in IaaS, including network controls. You share responsibility with us for application-level controls and IAM management. You remain responsible for your data security and client protection.
> https://cloud.google.com/architecture/framework/security/shared-responsibility-shared-fate#defined_by_workloads
>
> upvoted 2 times
>
> > **gcpengineer** 1 year, 3 months ago
> >
> > IaaS need more controls thn PaaS
> >
> > upvoted 1 times

tifo16 1 year, 8 months ago

Data at rest is encrypted by default by Google. So D is wrong. As mentioned by your link it Should be B.

upvoted 1 times

[Removed] 1 year, 12 months ago

Selected Answer: B

B it is.

upvoted 3 times

Question #80 *Topic 1*

An engineering team is launching a web application that will be public on the internet. The web application is hosted in multiple GCP regions and will be directed to the respective backend based on the URL request.
Your team wants to avoid exposing the application directly on the internet and wants to deny traffic from a specific list of malicious IP addresses.
Which solution should your team implement to meet these requirements?

A. Cloud Armor

B. Network Load Balancing

C. SSL Proxy Load Balancing

D. NAT Gateway

Correct Answer: *A*
Reference:
https://cloud.google.com/armor/docs/security-policy-concepts

DebasishLowes Highly Voted 👍 3 years, 5 months ago
Ans : A
upvoted 8 times

> **BillBaits** 2 years, 10 months ago
> Think so
> upvoted 1 times

Appsec977 Most Recent ⊙ 1 year, 3 months ago
Selected Answer: A
We can block the specific IPs in Cloud armor using simple rules or can use advanced rules using Common Expression Language(CEL).
upvoted 4 times

shayke 1 year, 8 months ago
Selected Answer: A
A Is the only ans because you are asked to limit access by IP and CA is the only option
upvoted 2 times

AzureDP900 1 year, 10 months ago
This is straight forward question, A is right
upvoted 1 times

AwesomeGCP 1 year, 11 months ago
Selected Answer: A
A. Cloud Armor
upvoted 2 times

cloudprincipal 2 years, 3 months ago
Selected Answer: A
https://cloud.google.com/armor/docs/security-policy-overview#edge-security
upvoted 2 times

[Removed] 3 years, 10 months ago
Ans - A
upvoted 4 times

mlyu 3 years, 11 months ago
Definitly B
upvoted 2 times

> **ownez** 3 years, 11 months ago
> Should be A? Cloud armor can deny traffic by defining IP addresses list rule and to avoid exposing the application directly on the internet.
>
> While Network LB is using Google Cloud firewalls to control or filter access to the backend VMs.
>
> Answer is A.
> upvoted 5 times
>
> > **mlyu** 3 years, 10 months ago

you are correct. Answer is A
The Cloud armor able to directed user traffic to an external HTTP(S) load balancer enters the PoP closest to the user in Premium Tier.
https://cloud.google.com/armor/docs/security-policy-overview#edge-security

upvoted 5 times

Question #81 *Topic 1*

A customer is running an analytics workload on Google Cloud Platform (GCP) where Compute Engine instances are accessing data stored on Cloud Storage.
Your team wants to make sure that this workload will not be able to access, or be accessed from, the internet.
Which two strategies should your team use to meet these requirements? (Choose two.)

A. Configure Private Google Access on the Compute Engine subnet

B. Avoid assigning public IP addresses to the Compute Engine cluster.

C. Make sure that the Compute Engine cluster is running on a separate subnet.

D. Turn off IP forwarding on the Compute Engine instances in the cluster.

E. Configure a Cloud NAT gateway.

Correct Answer: *AB*

MohitA Highly Voted 4 years ago

AB suits well

upvoted 20 times

DebasishLowes Highly Voted 3 years, 5 months ago

Ans : AB

upvoted 7 times

[Removed] Most Recent 1 year, 1 month ago

Selected Answer: AB

A,B
Has to be A and B together. A (Private Google Access) has minimal effect on instances with public IP so we also need to avoid assigning public IP to get the desired (internal only) effect.

https://cloud.google.com/vpc/docs/private-google-access

upvoted 2 times

gcpengineer 1 year, 3 months ago

Selected Answer: AB

AB, A to access the cloud storage privately

upvoted 2 times

gcpengineer 1 year, 3 months ago

Selected Answer: BE

BE. no public ip in vm and nat to access the cloud storage

upvoted 1 times

> **gcpengineer** 1 year, 3 months ago
>
> AB, A to access the cloud storage privately
>
> upvoted 1 times

therealsohail 1 year, 7 months ago

AE
A. Configuring Private Google Access on the Compute Engine subnet: This feature enables instances without public IP addresses to connect to Google APIs and services over internal IP addresses, ensuring that the instances cannot be accessed from the internet.

E. Configuring a Cloud NAT gateway: This ensures that instances within the VPC can connect to the internet, but only to specific IP ranges and po and it also ensures that the instances cannot initiate connection to the internet.

By configuring both options, you are providing your Compute Engine instances with a way to access Google services while also being isolated fro the internet and that is the best way to ensure that this workload will not be able to access, or be accessed from, the internet.

upvoted 1 times

> **diasporabro** 1 year, 7 months ago
>
> NAT Gateway allows an instance to access the public internet (while not being accessible from the public internet), so it is incorrect
>
> upvoted 2 times

AzureDP900 1 year, 10 months ago

AB is correct
A. Configure Private Google Access on the Compute Engine subnet
B. Avoid assigning public IP addresses to the Compute Engine cluster.

upvoted 1 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: AB

A. Configure Private Google Access on the Compute Engine subnet
B. Avoid assigning public IP addresses to the Compute Engine cluster.

upvoted 2 times

cloudprincipal 2 years, 3 months ago

Selected Answer: AB

agree with all the others

upvoted 2 times

pfilourenco 3 years, 3 months ago

B and E:
"make sure that this workload will not be able to access, or be accessed from, the internet."
If we have cloud NAT we are able to access the internet! Also with public IP.

upvoted 2 times

Rupo7 2 years, 6 months ago

The question says " not be able to access, or be accessed from, the internet." A NAT gateway enables access to the internet, just behind a stati IP. A. Private access for the subnet is required to enable access to GCS. B is a good measure, as then the instance cannot access the internet at all (without a NAT Gateway that is).

upvoted 1 times

gcpengineer 1 year, 3 months ago

private access of storage is required not of the VMs

upvoted 1 times

[Removed] 3 years, 4 months ago

Not A https://cloud.google.com/vpc/docs/private-google-access

upvoted 1 times

[Removed] 3 years, 4 months ago

NOt D, because by de fault IP forwarding is disabled. You do not need to turn it off.

upvoted 1 times

[Removed] 3 years, 4 months ago

So B and E is the right answer.

upvoted 3 times

tanfromvn 3 years, 2 months ago

A_B, why not A? Private access just accepts traffic in GCP and to GG API

upvoted 2 times

ffdd1234 3 years, 7 months ago

if you Avoid assigning public IP addresses to the Compute Engine cluster the instance could access to internet if have a nat gateway, maybe the answer is A and D

upvoted 1 times

ffdd1234 2 years, 10 months ago

+1 A-D

upvoted 1 times

ffdd1234 2 years, 10 months ago

But not sure "Ensure that IP Forwarding feature is not enabled at the Google Compute Engine instance level for security and compliance reasons, as instances with IP Forwarding enabled act as routers/packet forwarders."
IP FW is for route packets could not be D

upvoted 1 times

Topsy 3 years, 8 months ago

A and B is correct

upvoted 4 times

[Removed] 3 years, 10 months ago

Ans - AB

upvoted 2 times

genesis3k 3 years, 10 months ago

AB is the correct answer.

upvoted 1 times

Wooky 3 years, 11 months ago

B,D not A
Private google access provides public google api access without public IP

upvoted 1 times

Wooky 3 years, 11 months ago

My mistake, ans is AB.

upvoted 2 times

Raushanr 3 years, 11 months ago

Answer-AB

upvoted 2 times

Question #82 *Topic 1*

A customer wants to run a batch processing system on VMs and store the output files in a Cloud Storage bucket. The networking and security teams have decided that no VMs may reach the public internet.
How should this be accomplished?

A. Create a firewall rule to block internet traffic from the VM.

B. Provision a NAT Gateway to access the Cloud Storage API endpoint.

C. Enable Private Google Access.

D. Mount a Cloud Storage bucket as a local filesystem on every VM.

Correct Answer: *B*

tanfromvn Highly Voted 3 years, 2 months ago
C-there is no traffic to outside internet
upvoted 15 times

mynk29 2 years, 6 months ago
Private google access is enabled at subnet level not at VPC level.
upvoted 1 times

nilopo Highly Voted 1 year, 11 months ago
Selected Answer: C
The ask is to store the output files in a Cloud storage bucket. "The networking and security teams have decided that no VMs may reach the public internet" - No VMs MAY reach public internet but not 'MUST'. Hence 'C' is the answer
upvoted 7 times

desertlotus1211 Most Recent 6 months, 3 weeks ago
What if the VM is on-premise? The question never said it was in GCP?

Would the answer not be 'B'?
upvoted 1 times

Portugapt 7 months, 2 weeks ago
Selected Answer: C
What should be accomplished is the access to GCS, knowing VMs cannot access the public network.
So, Private Google Access accomplishes it.
upvoted 1 times

desertlotus1211 8 months ago
The answer is A....
With GPA enabled, VMs can still reach the Internet. Accessing the backend storage is ther to throw you off of what is being asked - and that's NO VMs may reach the Internet...

Answer is A
upvoted 1 times

[Removed] 8 months, 3 weeks ago
Selected Answer: C
C private google access allows access to google services without internet connection
upvoted 2 times

Xoxoo 11 months, 2 weeks ago
Selected Answer: C
To ensure that VMs can access Cloud Storage without reaching the public internet, you should:

C. Enable Private Google Access.

Enabling Private Google Access allows VMs with only internal IP addresses in a VPC network to access Google Cloud services like Cloud Storage without needing external IP addresses or going through the public internet.
upvoted 2 times

Xoxoo 11 months, 2 weeks ago
Option B, provisioning a NAT Gateway, would enable VMs to access the public internet, which is not in line with the requirement of not allowir VMs to reach the public internet.

Options A and D are not suitable for the specific requirement of accessing Cloud Storage while preventing VMs from reaching the public internet.

upvoted 1 times

blacortik 1 year ago

Selected Answer: B

B. Provision a NAT Gateway to access the Cloud Storage API endpoint.

Explanation:

To ensure that VMs can't reach the public internet but can still access Google Cloud services like Cloud Storage, you can use a Network Address Translation (NAT) Gateway. NAT Gateway allows instances in a private subnet to initiate outbound connections to the internet while masking their actual internal IP addresses. This way, the VMs can access the Cloud Storage API endpoint without directly connecting to the public internet.

upvoted 2 times

[Removed] 1 year, 1 month ago

Selected Answer: C

"C"
The question is not worded well. If you replace "..has decided.." with "..has enforced.." then the meat of the question becomes how to achieve the first part of the requirement which is reaching cloud storage without public access, which is through private google access.
Reference:
https://cloud.google.com/vpc/docs/private-google-access

upvoted 3 times

desertlotus1211 1 year ago

This has no effect and is meaningless if the VM has an external IP... You need to read the document:
'Private Google Access has no effect on instances that have external IP addresses. Instances with external IP addresses can access the internet, according to the internet access requirements'...

No where in the question say the VMs has or hasn't have an ext. IP.

Correct Answer is A

upvoted 1 times

gcpengineer 1 year, 3 months ago

Selected Answer: A

I think A is correct

upvoted 1 times

gcpengineer 1 year, 3 months ago

Selected Answer: B

B is the ans, as nat is needed to reach the cloud storage

upvoted 1 times

gcpengineer 1 year, 3 months ago

I think A is correct

upvoted 1 times

Lyfedge 1 year, 5 months ago

The question says "The networking and security teams have decided that no VMs may reach the public internet"y
A

upvoted 1 times

gcpengineer 1 year, 3 months ago

How are u suppose to access cloud storage?

upvoted 1 times

desertlotus1211 8 months ago

that not what they asked... they asked 'The networking and security teams have decided that no VMs may reach the public internet'.... so what do you do?

upvoted 1 times

Meyucho 1 year, 8 months ago

C!!!! This example is just the exact and only meaning for have PGA!!!

upvoted 1 times

TonytheTiger 1 year, 9 months ago

Answer C:
Here is why; the VM need to access google service i.e. "Cloud Storage Bucket".
Google doc states: Private Google Access permits access to Google APIs and services in Google's production infrastructure
https://cloud.google.com/vpc/docs/private-google-access
Everyone is reading the question as limited access to public internet but is missing the 2nd part of the question, which is access a google services ONLY enable Private Google Access will fulfil the requirement.

upvoted 2 times

Littleivy 1 year, 9 months ago

Selected Answer: C

C is the answer

upvoted 1 times

rotorclear 1 year, 10 months ago

Selected Answer: C

The ask is to access cloud storage while doing the batch processing not how to block the internet.
Overall it's a poor choice of words in the question attempting to confuse than check knowledge

upvoted 1 times

AzureDP900 1 year, 10 months ago

C is right

upvoted 1 times

AwesomeGCP 1 year, 11 months ago

C. Enable Private Google Access on the VPC.

upvoted 1 times

Question #83 *Topic 1*

As adoption of the Cloud Data Loss Prevention (Cloud DLP) API grows within your company, you need to optimize usage to reduce cost. Cloud DLP target data is stored in Cloud Storage and BigQuery. The location and region are identified as a suffix in the resource name.
Which cost reduction options should you recommend?

A. Set appropriate rowsLimit value on BigQuery data hosted outside the US and set appropriate bytesLimitPerFile value on multiregional Cloud Storage buckets.

B. Set appropriate rowsLimit value on BigQuery data hosted outside the US, and minimize transformation units on multiregional Cloud Storage buckets.

C. Use rowsLimit and bytesLimitPerFile to sample data and use CloudStorageRegexFileSet to limit scans.

D. Use FindingLimits and TimespanContfig to sample data and minimize transformation units.

Correct Answer: *C*
Reference:
https://cloud.google.com/dlp/docs/reference/rest/v2/InspectJobConfig

[Removed] Highly Voted 3 years, 10 months ago

Ans - C
https://cloud.google.com/dlp/docs/inspecting-storage#sampling
https://cloud.google.com/dlp/docs/best-practices-costs#limit_scans_of_files_in_to_only_relevant_files

upvoted 14 times

> **[Removed]** 3 years, 10 months ago
>
> https://cloud.google.com/dlp/docs/inspecting-storage#limiting-gcs
>
> upvoted 1 times

passtest100 Highly Voted 3 years, 11 months ago

C is the right one.

upvoted 5 times

Xoxoo Most Recent 11 months, 3 weeks ago

Selected Answer: C

To optimize usage of the Cloud Data Loss Prevention (Cloud DLP) API and reduce cost, you should consider using sampling and CloudStorageRegexFileSet to limit scans 1.

By sampling data, you can limit the amount of data that the DLP API scans, thereby reducing costs 1. You can use the rowsLimit and bytesLimitPerFile options to sample data and limit scans to specific files in Cloud Storage 1. You can also use CloudStorageRegexFileSet to limit scans to only specific files in Cloud Storage 1.

In addition, you can set appropriate rowsLimit value on BigQuery data hosted outside the US to further optimize usage and reduce costs 1.

upvoted 2 times

AzureDP900 1 year, 10 months ago

C is right

upvoted 4 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: C

C . Use rowsLimit and bytesLimitPerFile to sample data and use CloudStorageRegexFileSet to limit scans.

upvoted 4 times

cloudprincipal 2 years, 3 months ago

Selected Answer: C

https://cloud.google.com/dlp/docs/inspecting-storage#sampling

upvoted 3 times

Question #84 *Topic 1*

Your team uses a service account to authenticate data transfers from a given Compute Engine virtual machine instance of to a specified Cloud Storage bucket. An engineer accidentally deletes the service account, which breaks application functionality. You want to recover the application as quickly as possible without compromising security.
What should you do?

A. Temporarily disable authentication on the Cloud Storage bucket.

B. Use the undelete command to recover the deleted service account.

C. Create a new service account with the same name as the deleted service account.

D. Update the permissions of another existing service account and supply those credentials to the applications.

Correct Answer: *B*
Reference:
https://cloud.google.com/iam/docs/creating-managing-service-accounts#undeleting_a_service_account

DebasishLowes Highly Voted 3 years, 5 months ago
Ans : B
upvoted 9 times

saurabh1805 Highly Voted 3 years, 10 months ago
B is correct answer here.

https://cloud.google.com/iam/docs/reference/rest/v1/projects.serviceAccounts/undelete
upvoted 7 times

AzureDP900 1 year, 10 months ago
Thank you for sharing link, I agree B is right
upvoted 1 times

pradoUA Most Recent 11 months, 1 week ago
Selected Answer: B
B is correct
upvoted 2 times

ArizonaClassics 11 months, 3 weeks ago
B. Use the undelete command to recover the deleted service account.

Google Cloud Platform provides an undelete command that can be used to recover a recently deleted service account. This would be the fastest and most direct way to restore functionality without compromising security or introducing changes to the application configuration.
upvoted 3 times

[Removed] 1 year, 1 month ago
Selected Answer: B
"B"
Answer is B however the documentation has been updated. Not all links in other comments are valid still. Here's the latest link around this topic.
https://cloud.google.com/iam/docs/service-accounts-delete-undelete#undeleting
upvoted 3 times

AwesomeGCP 1 year, 11 months ago
Selected Answer: B
B. Use the undelete command to recover the deleted service account.
upvoted 3 times

[Removed] 3 years, 10 months ago
Ans - B
upvoted 3 times

MohitA 4 years ago
B is the Answer
upvoted 4 times

Question #85 *Topic 1*

You are the Security Admin in your company. You want to synchronize all security groups that have an email address from your LDAP directory in Cloud IAM.
What should you do?

A. Configure Google Cloud Directory Sync to sync security groups using LDAP search rules that have ג€user email addressג€ as the attribute to facilitate one-way sync.

B. Configure Google Cloud Directory Sync to sync security groups using LDAP search rules that have ג€user email addressג€ as the attribute to facilitate bidirectional sync.

C. Use a management tool to sync the subset based on the email address attribute. Create a group in the Google domain. A group created in a Google domain will automatically have an explicit Google Cloud Identity and Access Management (IAM) role.

D. Use a management tool to sync the subset based on group object class attribute. Create a group in the Google domain. A group created in a Google domain will automatically have an explicit Google Cloud Identity and Access Management (IAM) role.

Correct Answer: *A*

sudarchary Highly Voted 2 years, 7 months ago
Selected Answer: A
search rules that have "user email address" as the attribute to facilitate one-way sync.
Reference Links:
https://support.google.com/a/answer/6126589?hl=en
upvoted 11 times

JoseMaria111 Highly Voted 1 year, 11 months ago
GCDS allow sync ldap users in one way. A is correct
upvoted 5 times

GCBC Most Recent 1 year ago
A is correct
upvoted 2 times

PST21 1 year, 8 months ago
A is correct as it shoud be one way sync - LDAP -> Cloud Identity via GCDS
upvoted 2 times

AzureDP900 1 year, 10 months ago
A is correct
upvoted 3 times

AwesomeGCP 1 year, 11 months ago
Selected Answer: A
A. Configure Google Cloud Directory Sync to sync security groups using LDAP search rules that have “user email address” as the attribute to facilitate one-way sync.
upvoted 2 times

[Removed] 3 years, 4 months ago
Why A is not correct? GCP provide this sync tool.
upvoted 3 times

mistryminded 2 years, 9 months ago
Incorrect. GCDS is Google Workspace Admin tool.

Correct answer is A. GCDS only syncs one way - https://support.google.com/a/answer/106368?hl=en
upvoted 4 times

DebasishLowes 3 years, 5 months ago
Ans : A
upvoted 2 times

[Removed] 3 years, 10 months ago
Ans - A
upvoted 2 times

saurabh1805 3 years, 10 months ago
A is correct answer here.

upvoted 2 times

passtest100 3 years, 11 months ago

Answer - A

upvoted 2 times

skshak 3 years, 11 months ago

Answer - A

upvoted 2 times

Question #86 *Topic 1*

You are part of a security team investigating a compromised service account key. You need to audit which new resources were created by the service account.
What should you do?

A. Query Data Access logs.

B. Query Admin Activity logs.

C. Query Access Transparency logs.

D. Query Stackdriver Monitoring Workspace.

Correct Answer: *A*
Reference:
https://cloud.google.com/iam/docs/audit-logging/examples-service-accounts

MohitA Highly Voted 4 years ago
B is the Ans
upvoted 14 times

ownez 3 years, 11 months ago
Shouldn't it be A? The question is about which resources were created by the SA.

B (Admin Activity logs) cannot view this. It is only for user's activity such as create, modify or delete a particular SA.
upvoted 1 times

FatCharlie 3 years, 9 months ago
"Admin Activity audit logs contain log entries for API calls or other administrative actions that modify the configuration or metadata of resources. For example, these logs record when users create VM instances or change Identity and Access Management permissions".

This is exactly what you want to see. What resources were created by the SA?

https://cloud.google.com/logging/docs/audit#admin-activity
upvoted 10 times

AzureDP900 1 year, 10 months ago
B is right . Agree with your explanation
upvoted 2 times

Fellipo 3 years, 9 months ago
B it's OK
upvoted 4 times

VicF Highly Voted 3 years, 4 months ago
Ans B
"B" is for actions that modify the configuration or metadata of resources. For example, these logs record when users create VM instances or chan
Identity and Access Management permissions.
"A" is only for "user-provided" resource data. Data Access audit logs-- except for BigQuery Data Access audit logs-- "are disabled by default"
upvoted 6 times

dija123 Most Recent 5 months, 1 week ago
Selected Answer: B
Agree with B
upvoted 1 times

Xoxoo 11 months, 3 weeks ago
Selected Answer: B
To audit which new resources were created by a compromised service account key, you should query Admin Activity logs 1.

Admin Activity logs provide a record of every administrative action taken in your Google Cloud Platform (GCP) project, including the creation of new resources 1. By querying Admin Activity logs, you can identify which new resources were created by the compromised service account key ar
take appropriate action to secure your environment 1.

You can use the gcloud command-line tool or the Cloud Console to query Admin Activity logs 1. You can filter the logs based on specific criteria, such as time range, user, or resource type 1.
upvoted 2 times

Meyucho 1 year, 8 months ago

Selected Answer: B

B - Audit logs. They have all the API calls that creates, modify or destroy resources. https://cloud.google.com/logging/docs/audit#admin-activity

upvoted 2 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: B

B. Query Admin Activity logs.

upvoted 3 times

JoseMaria111 1 year, 11 months ago

Admin activity log records resources changes. B is correct

upvoted 2 times

piyush_1982 2 years, 1 month ago

Selected Answer: B

Admin activity logs are always created to log entries for API calls or other actions that modify the configuration or metadata of resources. For example, these logs record when users create VM instances or change Identity and Access Management permissions.

upvoted 2 times

cloudprincipal 2 years, 3 months ago

Selected Answer: B

Admin activity logs contain all GCP API calls.
So this is where the service account activity will show up

upvoted 2 times

[Removed] 3 years, 4 months ago

I support B, https://cloud.google.com/iam/docs/audit-logging
says IAM logs write into admin log

upvoted 4 times

DebasishLowes 3 years, 5 months ago

Ans : B

upvoted 3 times

[Removed] 3 years, 10 months ago

Ans - B

upvoted 4 times

Question #87 *Topic 1*

You have an application where the frontend is deployed on a managed instance group in subnet A and the data layer is stored on a mysql Compute Engine virtual machine (VM) in subnet B on the same VPC. Subnet A and Subnet B hold several other Compute Engine VMs. You only want to allow the application frontend to access the data in the application's mysql instance on port 3306.
What should you do?

A. Configure an ingress firewall rule that allows communication from the src IP range of subnet A to the tag "data-tag" that is applied to the mysql Compute Engine VM on port 3306.

B. Configure an ingress firewall rule that allows communication from the frontend's unique service account to the unique service account of the mysql Compute Engine VM on port 3306.

C. Configure a network tag "fe-tag" to be applied to all instances in subnet A and a network tag "data-tag" to be applied to all instances in subnet B. Then configure an egress firewall rule that allows communication from Compute Engine VMs tagged with data-tag to destination Compute Engine VMs tagged fe- tag.

D. Configure a network tag "fe-tag" to be applied to all instances in subnet A and a network tag "data-tag" to be applied to all instances in subnet B. Then configure an ingress firewall rule that allows communication from Compute Engine VMs tagged with fe-tag to destination Compute Engine VMs tagged with data-tag.

Correct Answer: *B*

Zuy01 Highly Voted 👍 3 years ago

B for sure, u can check this :
https://cloud.google.com/sql/docs/mysql/sql-proxy#using-a-service-account

upvoted 11 times

dija123 Most Recent 🕓 5 months, 2 weeks ago

Selected Answer: B

Agree with B

upvoted 1 times

Xoxoo 11 months, 3 weeks ago

Selected Answer: B

This approach ensures that only the application frontend can access the data in the MySQL instance, while all other Compute Engine VMs in subn A and subnet B are restricted from accessing it .

By configuring an ingress firewall rule that allows communication between the frontend's unique service account and the unique service account the MySQL Compute Engine VM, you can ensure that only authorized users can access your MySQL instance .

upvoted 2 times

GCBC 1 year ago

B Firellas rules using service account is better than tag

upvoted 2 times

[Removed] 1 year, 1 month ago

Selected Answer: B

"B"
I believe the answer is between B and A since part of the requirement is specifying the port. B is more correct since it leverages service accounts which is best practice for authentication/communication between application and database. Also, answer "A" allows ALL instances in the subnet t reach to reach mysql which is not desired. They only want the specific Frontend instances to reach excluding other instances in the subnet.

https://cloud.google.com/firewall/docs/firewalls#best_practices_for_firewall_rules

upvoted 3 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: B

B. Configure an ingress firewall rule that allows communication from the frontend's unique service account to the unique service account of the mysql ComputeEngine VM on port 3306.

upvoted 3 times

JoseMaria111 1 year, 11 months ago

B is correct.firellas rules using service account is better than tag based.
https://cloud.google.com/vpc/docs/firewalls#best_practices_for_firewall_rules

upvoted 2 times

mT3 2 years, 3 months ago

Selected Answer: B

Ans : B

upvoted 4 times

major_querty 2 years, 9 months ago

why is it not a?
a seems straight forward

The link which Zuy01 provided for answer b states: For this reason, using a service account is the recommended method for production instances NOT running on a Compute Engine instance.

upvoted 4 times

Arturo_Cloud 2 years ago

I agree (A), it is planned to limit a MySQL server in Compute Engine (IaaS) not in Cloud SQL (PaaS), so Networks Tags is the most common and recommended to use. Don't get confused with the services....

upvoted 2 times

Loved 1 year, 9 months ago

But answer A says "communication from the src IP range of subnet A"... this rules include all the instances on subnet A, while you have to consider only the frontend

upvoted 1 times

DebasishLowes 3 years, 5 months ago

Ans : B

upvoted 2 times

dtmtor 3 years, 5 months ago

ans is B

upvoted 2 times

[Removed] 3 years, 10 months ago

Ans - B

upvoted 4 times

Rantu 3 years, 11 months ago

B is correct

upvoted 4 times

Question #88 *Topic 1*

Your company operates an application instance group that is currently deployed behind a Google Cloud load balancer in us-central-1 and is configured to use the
Standard Tier network. The infrastructure team wants to expand to a second Google Cloud region, us-east-2. You need to set up a single external IP address to distribute new requests to the instance groups in both regions.
What should you do?

A. Change the load balancer backend configuration to use network endpoint groups instead of instance groups.

B. Change the load balancer frontend configuration to use the Premium Tier network, and add the new instance group.

C. Create a new load balancer in us-east-2 using the Standard Tier network, and assign a static external IP address.

D. Create a Cloud VPN connection between the two regions, and enable Google Private Access.

Correct Answer: *A*

Fellipo Highly Voted 3 years, 9 months ago

In Premium Tier: Backends can be in any region and any VPC network.

In Standard Tier: Backends must be in the same region as the forwarding rule, but can be in any VPC network.

upvoted 14 times

AzureDP900 1 year, 10 months ago

B is right

upvoted 2 times

mlyu Highly Voted 3 years, 11 months ago

Should be B
In Standard Tier LB, Backends must be in the same region
https://cloud.google.com/load-balancing/docs/load-balancing-overview#backend_region_and_network

upvoted 8 times

hakunamatataa Most Recent 11 months, 2 weeks ago

Selected Answer: B

B is the correct answer.

upvoted 2 times

Xoxoo 11 months, 3 weeks ago

Selected Answer: B

To set up a single external IP address to distribute new requests to the instance groups in both regions, you should change the load balancer frontend configuration to use the Premium Tier network, and add the new instance group .

By changing the load balancer frontend configuration to use the Premium Tier network, you can create a global load balancer that can distribute traffic across multiple regions using a single IP address . You can then add the new instance group to the existing load balancer to ensure that ne requests are distributed to both regions .

This approach provides a scalable and cost-effective solution for distributing traffic across multiple regions while ensuring high availability and lo latency .

upvoted 3 times

[Removed] 1 year, 1 month ago

Selected Answer: B

"B"
Answer is "B". Premium Network Tier allows you to span multiple regions.

https://cloud.google.com/network-tiers

upvoted 4 times

bwrutu 1 year, 1 month ago

Can somebody mail all the 185 questions in the pdf format to wruthuraj@gmail.com. Iam not getting access after the 20th question set

upvoted 1 times

spoxman 1 year, 5 months ago

Selected Answer: B

only Premium allows LB between regions

upvoted 1 times

Meyucho 1 year, 8 months ago

Selected Answer: B

Global load balancers require Premium Tier!

upvoted 1 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: B

B. Change the load balancer frontend configuration to use the Premium Tier network, and add the new instance group.

upvoted 1 times

cloudprincipal 2 years, 3 months ago

Selected Answer: B

https://cloud.google.com/load-balancing/docs/choosing-load-balancer#global-regional

upvoted 1 times

DebasishLowes 3 years, 5 months ago

Ans : B

upvoted 2 times

saurabh1805 3 years, 10 months ago

I will also go with Option B

upvoted 6 times

Question #89 *Topic 1*

You are the security admin of your company. You have 3,000 objects in your Cloud Storage bucket. You do not want to manage access to each object individually.
You also do not want the uploader of an object to always have full control of the object. However, you want to use Cloud Audit Logs to manage access to your bucket.
What should you do?

A. Set up an ACL with OWNER permission to a scope of allUsers.

B. Set up an ACL with READER permission to a scope of allUsers.

C. Set up a default bucket ACL and manage access for users using IAM.

D. Set up Uniform bucket-level access on the Cloud Storage bucket and manage access for users using IAM.

Correct Answer: *A*
Reference:
https://cloud.google.com/storage/docs/access-control/lists

Fellipo Highly Voted 3 years, 9 months ago
it's D, https://cloud.google.com/storage/docs/uniform-bucket-level-access#:~:text=When%20you%20enable%20uniform%20bucket,and%20the%20objects%20it%20contains.
upvoted 19 times

AwesomeGCP Highly Voted 1 year, 11 months ago
Selected Answer: D
D. Set up Uniform bucket-level access on the Cloud Storage bucket and manage access for users using IAM.
upvoted 5 times

tia_gll Most Recent 5 months, 2 weeks ago
Selected Answer: D
ans is D
upvoted 1 times

nccdebug 6 months, 2 weeks ago
Ans: D. https://cloud.google.com/storage/docs/uniform-bucket-level-access
upvoted 1 times

Xoxoo 11 months, 3 weeks ago
Selected Answer: D
To manage access to your Cloud Storage bucket without having to manage access to each object individually, you should set up Uniform bucket-level access on the Cloud Storage bucket and manage access for users using IAM .

Uniform bucket-level access allows you to use Identity and Access Management (IAM) alone to manage permissions for all objects contained insi the bucket or groups of objects with common name prefixes . This approach simplifies access management and ensures that all objects in the bucket have the same level of access .

By using IAM, you can grant users specific permissions to access your Cloud Storage bucket, such as read, write, or delete permissions . You can also use Cloud Audit Logs to monitor and manage access to your bucket .

This approach provides a secure environment for your Cloud Storage bucket while ensuring that only authorized users can access it .
upvoted 4 times

AzureDP900 1 year, 10 months ago
D is right
upvoted 3 times

cloudprincipal 2 years, 3 months ago
Selected Answer: D
https://cloud.google.com/storage/docs/uniform-bucket-level-access#enabled
upvoted 3 times

ramravella 3 years, 2 months ago
Answer is A. Read the note below in the below URL

https://cloud.google.com/storage/docs/access-control/lists

Note: You cannot grant discrete permissions for reading or writing ACLs or other metadata. To allow someone to read and write ACLs, you must grant them OWNER permission.

upvoted 1 times

Zuy01 3 years ago

the question mention "do not want the uploader of an object to always have full control of the object" that's mean you shouldn't grant the owner permission, hence the best ans is D.

upvoted 3 times

[Removed] 3 years, 4 months ago

A grants Owner???too much for this.

upvoted 2 times

[Removed] 3 years, 10 months ago

Ans - D

upvoted 3 times

saurabh1805 3 years, 10 months ago

I will go with uniform level access and manage access via IAM,

Hence D.

upvoted 2 times

passtest100 3 years, 11 months ago

SHOULD BE D

upvoted 2 times

skshak 3 years, 11 months ago

Answer C https://cloud.google.com/storage/docs/access-control
Uniform (recommended): Uniform bucket-level access allows you to use Identity and Access Management (IAM) alone to manage permissions. IA applies permissions to all the objects contained inside the bucket or groups of objects with common name prefixes. IAM also allows you to use features that are not available when working with ACLs, such as IAM Conditions and Cloud Audit Logs.

upvoted 1 times

skshak 3 years, 11 months ago

Sorry, It is D. It was typo.

upvoted 3 times

mlyu 3 years, 11 months ago

the question stated they need cloud audit log for the GCS access, however uniform bucket-level access has restriction on the cloud audit log.
See https://cloud.google.com/storage/docs/uniform-bucket-level-access
The following restrictions apply when using uniform bucket-level access:
Cloud Logging and Cloud Audit Logs cannot export to buckets that have uniform bucket-level access enabled.

upvoted 1 times

FatCharlie 3 years, 9 months ago

They're not saying they want to export the logs to the bucket. They're just saying they want to "use Cloud Audit Logs to manage access to your bucket" (whatever that means).

upvoted 1 times

Question #90 *Topic 1*

You are the security admin of your company. Your development team creates multiple GCP projects under the "implementation" folder for several dev, staging, and production workloads. You want to prevent data exfiltration by malicious insiders or compromised code by setting up a security perimeter. However, you do not want to restrict communication between the projects.
What should you do?

A. Use a Shared VPC to enable communication between all projects, and use firewall rules to prevent data exfiltration.

B. Create access levels in Access Context Manager to prevent data exfiltration, and use a shared VPC for communication between projects.

C. Use an infrastructure-as-code software tool to set up a single service perimeter and to deploy a Cloud Function that monitors the "implementation" folder via Stackdriver and Cloud Pub/Sub. When the function notices that a new project is added to the folder, it executes Terraform to add the new project to the associated perimeter.

D. Use an infrastructure-as-code software tool to set up three different service perimeters for dev, staging, and prod and to deploy a Cloud Function that monitors the "implementation" folder via Stackdriver and Cloud Pub/Sub. When the function notices that a new project is added to the folder, it executes Terraform to add the new project to the respective perimeter.

Correct Answer: *B*

jonclem Highly Voted 3 years, 9 months ago

I'd also go with option B and here's why:
https://cloud.google.com/access-context-manager/docs/overview

Option A was a consideration until I came across this: https://cloud.google.com/security/data-loss-prevention/preventing-data-exfiltration

upvoted 16 times

dzhu Highly Voted 3 years ago

I think this is C. Communication between the project is necessary tied to VPC, but you need to include all projects under implementation folder in single VPCSC

upvoted 11 times

Bettoxicity Most Recent 5 months, 1 week ago

Selected Answer: D

Similitudes con la opción C:
Uso de IaC y Cloud Function: La opción D también utiliza una herramienta de IaC (Terraform) y una Cloud Function para automatizar la creación y gestión de los service perimeters.
Monitoreo con Stackdriver y Cloud Pub/Sub: Se utiliza Stackdriver y Cloud Pub/Sub para detectar la creación de nuevos proyectos.

Diferencias con la opción C:
Cantidad de service perimeters: La opción D crea tres service perimeters diferentes (dev, staging, prod), mientras que la opción C solo crea uno.
Asignación automática de proyectos: La función Cloud de la opción D asigna automáticamente los nuevos proyectos al perímetro de servicio correspondiente. En la opción C, la asignación de proyectos a los service perimeters se debe realizar manualmente.

upvoted 1 times

Sukon_Desknot 6 months, 4 weeks ago

Selected Answer: D

Using Access Context Manager service perimeters provides a security boundary to prevent data exfiltration.
Separate perimeters for dev, staging, prod provides appropriate isolation.
Shared VPC allows communication between projects within the perimeter.
The Cloud Function automaticaly adds new projects to the right perimeter via Terraform.
This meets all requirements - security perimeter to prevent data exfiltration, communication between projects, and automatic perimeter assignment for new projects.

upvoted 1 times

ssk119 1 year ago

just having vpc alone does not protect with data exfiltration. The correct answer is B

upvoted 1 times

desertlotus1211 1 year ago

you'd have to re-create the projects as a Host VPC... can't do that... too much work

upvoted 1 times

[Removed] 1 year, 1 month ago

Selected Answer: C

"C"
As others noted, VPC Service Controls are designed specifically to protect against the risks described in the question. Only one Service perimeter needed which excludes "D".

https://cloud.google.com/vpc-service-controls/docs/overview#benefits

upvoted 2 times

fad3r 1 year, 5 months ago

This question is very old. The answer is VPC Service controls.

Highly doubt this is still relevant.

upvoted 5 times

soltium 1 year, 10 months ago

Selected Answer: C

C. The keyword "prevent data exfiltration by malicious insiders or compromised code" is listed as the benefits of VPC service control
https://cloud.google.com/vpc-service-controls/docs/overview#benefits

Only C and D creates service perimeters, but D creates three and doesn't specify a bridge to connect those service perimeters so I choose C as the answer.

upvoted 4 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: C

C. Use an infrastructure-as-code software tool to set up a single service perimeter and to deploy a Cloud Function that monitors the "implementation" folder via Stackdriver and Cloud Pub/Sub. When the function notices that a new project is added to the folder, it executes Terraform to add the new project to the associated perimeter.

upvoted 1 times

cloudprincipal 2 years, 3 months ago

Selected Answer: C

eshtanaka is right: https://github.com/terraform-google-modules/terraform-google-vpc-service-controls/tree/master/examples/automatic_folde

upvoted 3 times

sudarchary 2 years, 7 months ago

Answer is A. Please focus on "security perimeter" and
"compromised code".

upvoted 1 times

eshtanaka 2 years, 10 months ago

Correct answer is C. See the description for "automatically secured folder" https://github.com/terraform-google-modules/terraform-google-vpc-service-controls/tree/master/examples/automatic_folder

upvoted 3 times

nilb94 3 years ago

Think it should be C. Access Context Manager docs say it is for ingress. Service Controls seems correct for exfiltration, and projects must be allow to communicate with each other so they need to be in a single service perimeter.

upvoted 3 times

desertlotus1211 3 years, 5 months ago

Answer is B:
https://cloud.google.com/access-context-manager/docs/overview

You need to read the question AND Answer carefully before selecting.
Answer A is in Answer B

upvoted 2 times

DebasishLowes 3 years, 5 months ago

Ans : A. To make the communication between different projects, shared vpc is required.

upvoted 1 times

HateMicrosoft 3 years, 5 months ago

The correct anwser is :B
Access Context Manager
https://cloud.google.com/access-context-manager/docs/overview

Preventing Data Exfiltration
https://cloud.google.com/security/data-loss-prevention/preventing-data-exfiltration

upvoted 3 times

> **gcpengineer** 1 year, 3 months ago
>
> called vpc service control now
>
> upvoted 1 times

deardeer 3 years, 7 months ago

D is the answer. This question is about sharing across perimeters. https://cloud.google.com/vpc-service-controls/docs/share-across-perimeters#service_perimeter_bridges

upvoted 3 times

Question #91 *Topic 1*

You need to provide a corporate user account in Google Cloud for each of your developers and operational staff who need direct access to GCP resources.
Corporate policy requires you to maintain the user identity in a third-party identity management provider and leverage single sign-on. You learn that a significant number of users are using their corporate domain email addresses for personal Google accounts, and you need to follow Google recommended practices to convert existing unmanaged users to managed accounts.
Which two actions should you take? (Choose two.)

A. Use Google Cloud Directory Sync to synchronize your local identity management system to Cloud Identity.

B. Use the Google Admin console to view which managed users are using a personal account for their recovery email.

C. Add users to your managed Google account and force users to change the email addresses associated with their personal accounts.

D. Use the Transfer Tool for Unmanaged Users (TTUU) to find users with conflicting accounts and ask them to transfer their personal Google accounts.

E. Send an email to all of your employees and ask those users with corporate email addresses for personal Google accounts to delete the personal accounts immediately.

Correct Answer: *BE*

VicF Highly Voted 3 years, 4 months ago
A&D.
A- Requires third-party IDp and wants to leverage single sign-on.
D- https://cloud.google.com/architecture/identity/migrating-consumer-accounts#initiating_a_transfer
"In addition to showing you all unmanaged accounts, the transfer tool for unmanaged users lets you initiate an account transfer by sending an account transfer request."
upvoted 17 times

skshak Highly Voted 3 years, 11 months ago
Is the answer is A,D
A - Requirement is third-party identity management provider and leverage single sign-on.
D - https://cloud.google.com/architecture/identity/assessing-existing-user-accounts (Use the transfer tool for unmanaged users to identify consumer accounts that use an email address that matches one of the domains you've added to Cloud Identity or G Suite.)
upvoted 8 times

dsafeqf Most Recent 11 months, 1 week ago
C, D are correct - https://cloud.google.com/architecture/identity/assessing-existing-user-accounts
upvoted 1 times

Littleivy 1 year, 9 months ago
Selected Answer: AD
A to sync IdP
D to transfer unmanaged accounts
upvoted 3 times

AzureDP900 1 year, 10 months ago
AD is right
upvoted 2 times

AwesomeGCP 1 year, 11 months ago
Selected Answer: AD
A. Use Google Cloud Directory Sync to synchronize your local identity management system to Cloud Identity.
D. Use the Transfer Tool for Unmanaged Users (TTUU) to find users with conflicting accounts and ask them to transfer their personal Google accounts.
upvoted 4 times

cloudprincipal 2 years, 3 months ago
Selected Answer: AD
see other comments
upvoted 3 times

sudarchary 2 years, 7 months ago
Answers are: A&C
https://cloud.google.com/architecture/identity/assessing-existing-user-accounts
upvoted 1 times

CloudTrip 3 years, 6 months ago

The keyword is here "convert" follow Google recommended practices to convert existing unmanaged users to managed accounts. So why sync unmanaged with Cloud Identity. I would prefer Answers C and D

upvoted 2 times

ThisisJohn 2 years, 8 months ago

But dont forget about "Corporate policy requires you to maintain the user identity in a third-party identity management provider".

I believe that makes it A and D

upvoted 1 times

mikelabs 3 years, 9 months ago

Answer is C,D. From GSuite Console you can do both.

upvoted 2 times

[Removed] 3 years, 10 months ago

Ans - AD

upvoted 4 times

[Removed] 3 years, 10 months ago

https://cloud.google.com/architecture/identity/migrating-consumer-accounts#initiating_a_transfer

upvoted 7 times

saurabh1805 3 years, 10 months ago

A, D is correct answer

upvoted 4 times

lordb 3 years, 11 months ago

https://cloud.google.com/architecture/identity/assessing-existing-user-accounts

upvoted 2 times

Question #92 *Topic 1*

You are on your company's development team. You noticed that your web application hosted in staging on GKE dynamically includes user data in web pages without first properly validating the inputted data. This could allow an attacker to execute gibberish commands and display arbitrary content in a victim user's browser in a production environment.
How should you prevent and fix this vulnerability?

A. Use Cloud IAP based on IP address or end-user device attributes to prevent and fix the vulnerability.

B. Set up an HTTPS load balancer, and then use Cloud Armor for the production environment to prevent the potential XSS attack.

C. Use Web Security Scanner to validate the usage of an outdated library in the code, and then use a secured version of the included library.

D. Use Web Security Scanner in staging to simulate an XSS injection attack, and then use a templating system that supports contextual auto-escaping.

Correct Answer: *D*
Reference:
https://cloud.google.com/security-scanner/docs/remediate-findings

sudarchary Highly Voted 2 years, 7 months ago
Selected Answer: D
Option D is correct as using web security scanner will allow to detect the vulnerability and templating system
upvoted 10 times

deardeer Highly Voted 3 years, 7 months ago
Answer is D. There is mention about simulating in Web Security Scanner. "Web Security Scanner cross-site scripting (XSS) injection testing *simulates* an injection attack by inserting a benign test string into user-editable fields and then performing various user actions."
https://cloud.google.com/security-command-center/docs/how-to-remediate-web-security-scanner-findings#xss
upvoted 7 times

ThisisJohn 2 years, 8 months ago
Agree. Also from your link

"There are various ways to fix this problem. The recommended fix is to escape all output and use a templating system that supports contextua auto-escaping."

So escaping is a way to fix the issue, which is required by the question
upvoted 1 times

AzureDP900 1 year, 10 months ago
Agree with D
upvoted 2 times

[Removed] Most Recent 1 year, 1 month ago
Selected Answer: D
"D"
Using Web Security Scanner in Security Command Center to find XSS vulnerabilities. This page explains recommended mitigation techniques suc as using contextual auto-escaping.

https://cloud.google.com/security-command-center/docs/how-to-remediate-web-security-scanner-findings#xss
upvoted 2 times

AwesomeGCP 1 year, 11 months ago
Selected Answer: D
D. Use Web Security Scanner in staging to simulate an XSS injection attack, and then use a templating system that supports contextual auto-escaping.
upvoted 2 times

tangac 1 year, 12 months ago
Selected Answer: D
clear D everything is explicated here : https://cloud.google.com/security-command-center/docs/how-to-remediate-web-security-scanner-finding
Web Security Scanner cross-site scripting (XSS) injection testing simulates an injection attack by inserting a benign test string into user-editable fields and then performing various user actions. Custom detectors observe the browser and DOM during this test to determine whether an injection was successful and assess its potential for exploitation.
There are various ways to fix this issue. The recommended fix is to escape all output and use a templating system that supports contextual auto-escaping.
upvoted 2 times

Lancyqusa 2 years, 8 months ago

It should be C because the web security scanner will identify the library known to contain the security issue as in the examples here - https://cloud.google.com/security-command-center/docs/how-to-use-web-security-scanner#example_findings .
Once the security issue is identified, the vulnerability can be fixed by a secure version of that library.

upvoted 1 times

DebasishLowes 3 years, 5 months ago

Ans : D

upvoted 2 times

pyc 3 years, 7 months ago

C,
D is wrong, as Security Scanner can't "simulate" anything. It's a scanner.
B is not right, as Armor can't do input data validation, it just deny/allow IP/CIDR.

upvoted 1 times

> **desertlotus1211** 3 years, 5 months ago
>
> Yes it can simulate... Read the documentation first...
>
> upvoted 3 times

KarVaid 3 years, 8 months ago

https://cloud.google.com/security-command-center/docs/concepts-web-security-scanner-overview

Security Scanner should be able to scan for XSS vulnerabilities as well. Option D is better.

upvoted 2 times

> **KarVaid** 3 years, 8 months ago
>
> Cloud armor can prevent the vulnerability but to fix it, you would need Security scanner.
>
> upvoted 1 times

Fellipo 3 years, 9 months ago

B , https://cloud.google.com/armor

upvoted 5 times

[Removed] 3 years, 10 months ago

Ans - D

upvoted 3 times

HectorLeon2099 3 years, 10 months ago

Answer is B. Web Security Scanner can look for XSS vulnerabilities but can't simulate XSS injection attack.
https://cloud.google.com/armor/docs/rule-tuning#cross-site_scripting_xss

upvoted 3 times

> **saurabh1805** 3 years, 10 months ago
>
> Agree B is correct answer here.
>
> upvoted 2 times

> **FatCharlie** 3 years, 9 months ago
>
> Web Security Scanner does appear to be able to simulate an XSS attack.
>
> "Web Security Scanner cross-site scripting (XSS) injection testing simulates an injection attack by inserting a benign test string into user-editable fields and then performing various user actions. Custom detectors observe the browser and DOM during this test to determine whether an injection was successful and assess its potential for exploitation."
>
> https://cloud.google.com/security-command-center/docs/how-to-remediate-web-security-scanner-findings#remediate-findings
>
> upvoted 4 times

Jerrard 3 years, 11 months ago

D. https://cloud.google.com/security-command-center/docs/concepts-web-security-scanner-overview

upvoted 4 times

Question #93 *Topic 1*

You are part of a security team that wants to ensure that a Cloud Storage bucket in Project A can only be readable from Project B. You also want to ensure that data in the Cloud Storage bucket cannot be accessed from or copied to Cloud Storage buckets outside the network, even if the user has the correct credentials.
What should you do?

A. Enable VPC Service Controls, create a perimeter with Project A and B, and include Cloud Storage service.

B. Enable Domain Restricted Sharing Organization Policy and Bucket Policy Only on the Cloud Storage bucket.

C. Enable Private Access in Project A and B networks with strict firewall rules to allow communication between the networks.

D. Enable VPC Peering between Project A and B networks with strict firewall rules to allow communication between the networks.

Correct Answer: *B*
Reference:
https://cloud.google.com/resource-manager/docs/organization-policy/restricting-domains

FatCharlie Highly Voted 3 years, 9 months ago
The answer is A. This is question is covered by an example given for VPC Service Perimeters

https://cloud.google.com/vpc-service-controls/docs/overview#isolate
upvoted 20 times

AzureDP900 1 year, 10 months ago
A is right
upvoted 2 times

[Removed] Most Recent 1 year, 1 month ago
Selected Answer: A
"A"
VPC Service controls were created for this type of use case.

https://cloud.google.com/vpc-service-controls/docs/overview#isolate
upvoted 2 times

alleinallein 1 year, 5 months ago
Why not D?
upvoted 1 times

shayke 1 year, 8 months ago
Selected Answer: A
A - a classic VPCSC question
upvoted 2 times

AwesomeGCP 1 year, 11 months ago
Selected Answer: A
A. Enable VPC Service Controls, create a perimeter with Project A and B, and include Cloud Storage service.
upvoted 3 times

cloudprincipal 2 years, 3 months ago
Selected Answer: A
https://cloud.google.com/vpc-service-controls/docs/overview#isolate
upvoted 2 times

nilb94 3 years ago
A - VPC Service Controls
upvoted 3 times

jeeet_ 3 years, 3 months ago
Answer is most positively A.
VPC service controls lets Security team create fine-grained Perimeter across projects within organization.
-> Security perimeter for API-Based services like Bigtable instances, Storage and Bigquery datasets.. are a kind of super powers for VPC Service control.
well in my test, I chose option B, but
Domain Restricted Organization policies are for limiting resource sharing based on domain.
so if you're out in internet, and have credentials you still can access resources based on your domain access level. So B option is wrong.
upvoted 2 times

HateMicrosoft 3 years, 5 months ago

The correct answer is: A
This is obtained by the VPC Service Controls by the perimeter setup.

Overview of VPC Service Controls
https://cloud.google.com/vpc-service-controls/docs/overview

upvoted 2 times

jonclem 3 years, 9 months ago

I would say option A is a better fit due to VPC Service Controls.

upvoted 3 times

jonclem 3 years, 9 months ago

I'd be inclined to agree, option B seems a better fit. Here's my reasoning behind it:
https://cloud.google.com/access-context-manager/docs/overview

upvoted 1 times

> **jonclem** 3 years, 9 months ago
>
> please ignore this comment, wrong question.
>
> upvoted 1 times

saurabh1805 3 years, 10 months ago

what is being asked is data exfiltration as well and which can be only achieved via VPC permiter and created a bridge between both project.

upvoted 1 times

> **Ducle** 3 years, 10 months ago
>
> A is better
>
> upvoted 2 times

[Removed] 3 years, 10 months ago

Ans - B

upvoted 1 times

Jerrard 3 years, 11 months ago

B. https://cloud.google.com/resource-manager/docs/organization-policy/restricting-domains

upvoted 1 times

Question #94

Topic 1

You are responsible for protecting highly sensitive data in BigQuery. Your operations teams need access to this data, but given privacy regulations, you want to ensure that they cannot read the sensitive fields such as email addresses and first names. These specific sensitive fields should only be available on a need-to- know basis to the Human Resources team. What should you do?

A. Perform data masking with the Cloud Data Loss Prevention API, and store that data in BigQuery for later use.

B. Perform data redaction with the Cloud Data Loss Prevention API, and store that data in BigQuery for later use.

C. Perform data inspection with the Cloud Data Loss Prevention API, and store that data in BigQuery for later use.

D. Perform tokenization for Pseudonymization with the Cloud Data Loss Prevention API, and store that data in BigQuery for later use.

Correct Answer: *D*

Reference:

https://towardsdatascience.com/bigquery-pii-and-cloud-data-loss-prevention-dlp-take-it-to-the-next-level-with-data-catalog-c47c31bcf677

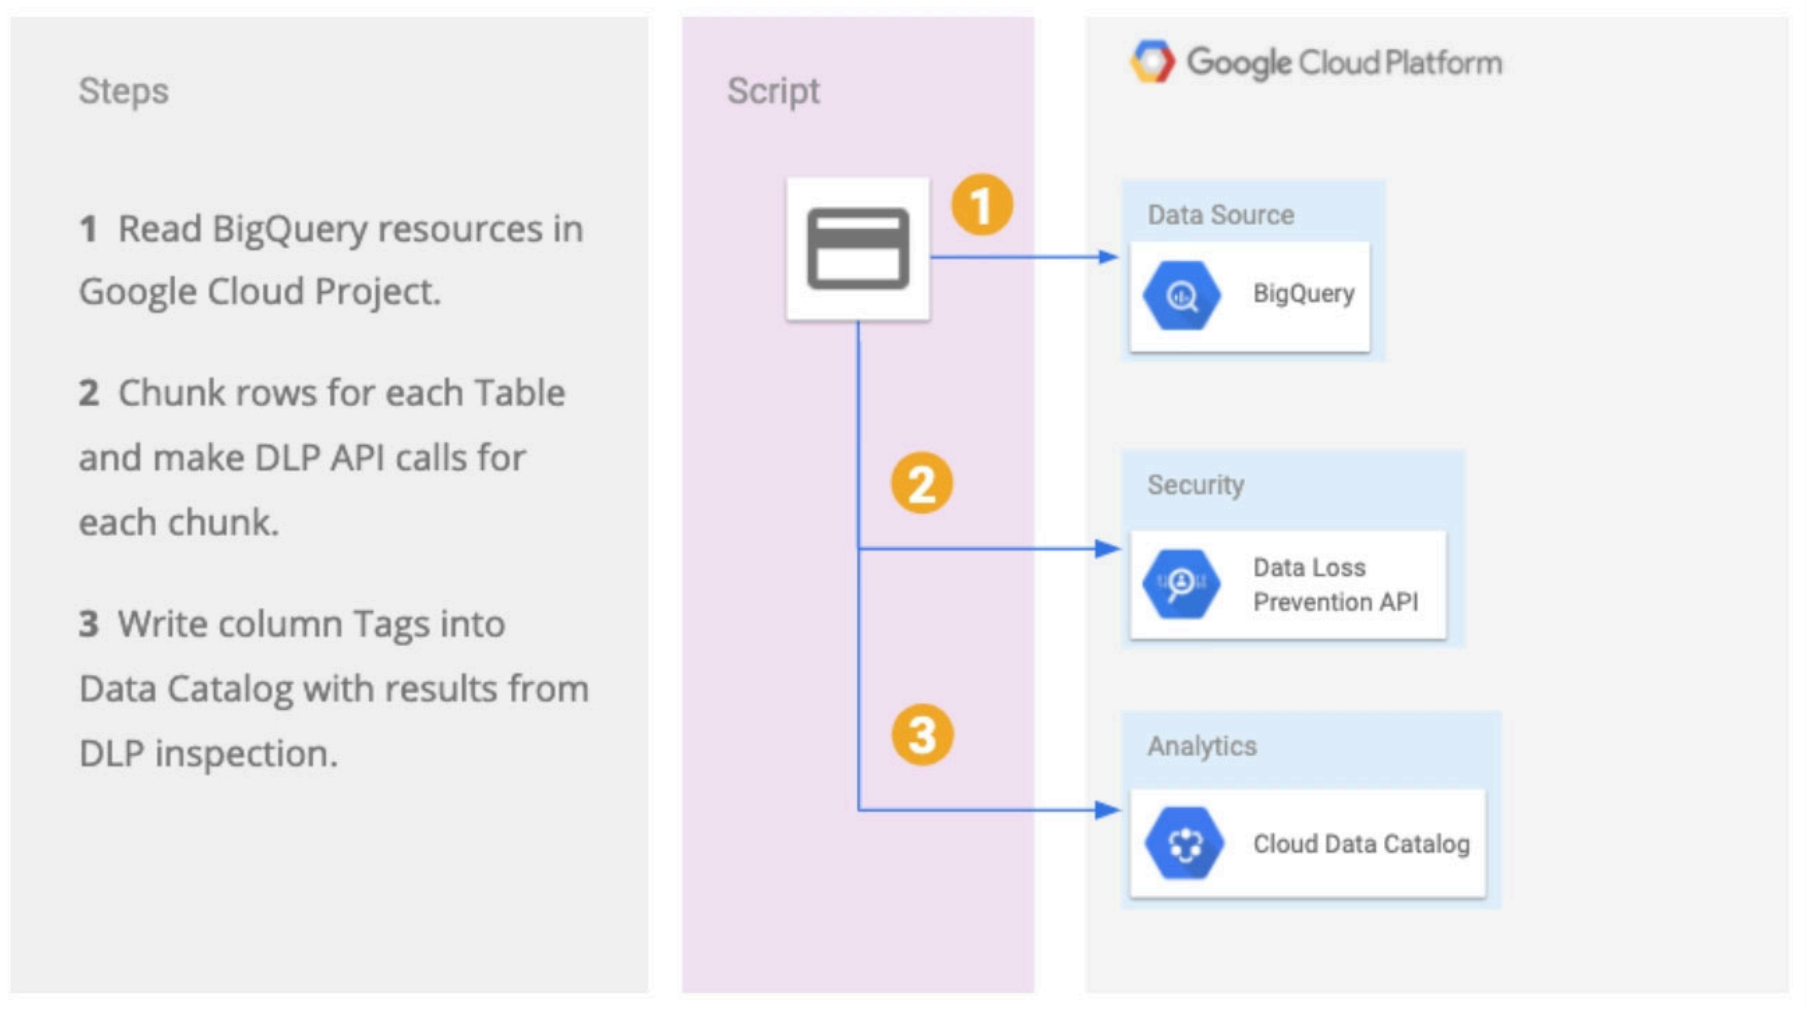

Steps

1 Read BigQuery resources in Google Cloud Project.

2 Chunk rows for each Table and make DLP API calls for each chunk.

3 Write column Tags into Data Catalog with results from DLP inspection.

zellck Highly Voted 1 year, 11 months ago

Selected Answer: D

D is the answer as tokenization can support re-identification for use by HR.

https://cloud.google.com/dlp/docs/pseudonymization

upvoted 5 times

AwesomeGCP Highly Voted 1 year, 11 months ago

Selected Answer: D

D. Perform tokenization for Pseudonymization with the Cloud Data Loss Prevention API, and store that data in BigQuery for later use.

upvoted 5 times

[Removed] Most Recent 1 year, 1 month ago

Selected Answer: D

"D"

Out of all the options listed, pseudonymization is the only reversible method which is one of the requirements in the quest.

https://cloud.google.com/dlp/docs/transformations-reference#transformation_methods
https://cloud.google.com/dlp/docs/pseudonymization

upvoted 3 times

Sammydp202020 1 year, 6 months ago

Selected Answer: D

Both A & D will do the job. But, A is preferred as the data is PII and needs to be secure.

https://cloud.google.com/dlp/docs/pseudonymization#how-tokenization-works

Why A is not a apt response:
https://cloud.google.com/bigquery/docs/column-data-masking-intro
The SHA-256 function used in data masking is type preserving, so the hash value it returns has the same data type as the column value.

SHA-256 is a deterministic hashing function; an initial value always resolves to the same hash value. However, it does not require encryption keys. This makes it possible for a malicious actor to use a brute force attack to determine the original value, by running all possible original values through the SHA-256 algorithm and seeing which one produces a hash that matches the hash returned by data masking.

upvoted 1 times

pedrojorge 1 year, 7 months ago

Selected Answer: D

D, as tokenization supports re-identification for the HR team

upvoted 2 times

therealsohail 1 year, 7 months ago

B is okay
Data redaction, as opposed to data masking or tokenization, completely removes or replaces the sensitive fields, making it so that the operations teams cannot see the sensitive information. This ensures that the sensitive data is only available to the Human Resources team on a need-to-knov basis, as per the privacy regulations. The Cloud Data Loss Prevention API is able to inspect and redact data, making it a suitable choice for this tas

upvoted 2 times

AzureDP900 1 year, 10 months ago

D is correct
Pseudonymization is a de-identification technique that replaces sensitive data values with cryptographically generated tokens. Pseudonymization widely used in industries like finance and healthcare to help reduce the risk of data in use, narrow compliance scope, and minimize the exposure sensitive data to systems while preserving data utility and accuracy.

upvoted 4 times

Random_Mane 1 year, 11 months ago

Selected Answer: A

A https://cloud.google.com/bigquery/docs/column-data-masking-intro

upvoted 3 times

heftjustice 1 year, 8 months ago

Data masking doesn't need DLP.

upvoted 2 times

Question #95 *Topic 1*

You are a Security Administrator at your organization. You need to restrict service account creation capability within production environments. You want to accomplish this centrally across the organization. What should you do?

A. Use Identity and Access Management (IAM) to restrict access of all users and service accounts that have access to the production environment.

B. Use organization policy constraints/iam.disableServiceAccountKeyCreation boolean to disable the creation of new service accounts.

C. Use organization policy constraints/iam.disableServiceAccountKeyUpload boolean to disable the creation of new service accounts.

D. Use organization policy constraints/iam.disableServiceAccountCreation boolean to disable the creation of new service accounts.

Correct Answer: *D*

Reference:

https://cloud.google.com/resource-manager/docs/organization-policy/restricting-service-accounts

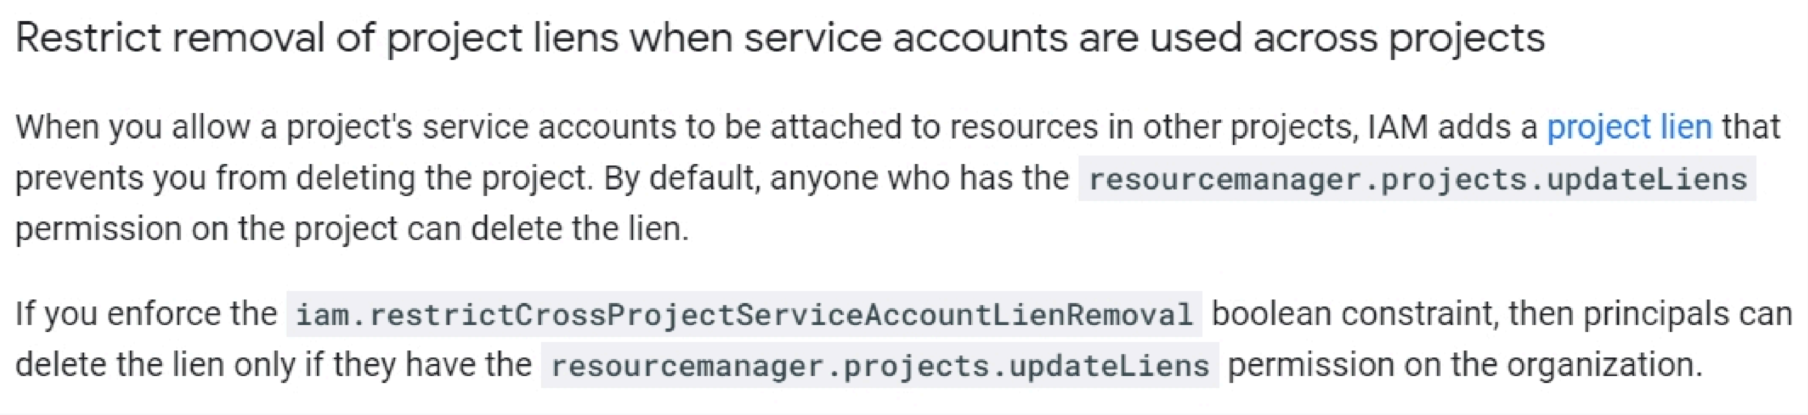

Restrict removal of project liens when service accounts are used across projects

When you allow a project's service accounts to be attached to resources in other projects, IAM adds a project lien that prevents you from deleting the project. By default, anyone who has the `resourcemanager.projects.updateLiens` permission on the project can delete the lien.

If you enforce the `iam.restrictCrossProjectServiceAccountLienRemoval` boolean constraint, then principals can delete the lien only if they have the `resourcemanager.projects.updateLiens` permission on the organization.

Tabayashi Highly Voted 2 years, 4 months ago

Answer is (D).

You can use the iam.disableServiceAccountCreation boolean constraint to disable the creation of new service accounts. This allows you to centrali
management of service accounts while not restricting the other permissions your developers have on projects.
https://cloud.google.com/resource-manager/docs/organization-policy/restricting-service-accounts#disable_service_account_creation

upvoted 11 times

[Removed] Highly Voted 1 year, 1 month ago

Selected Answer: D

"D"
Refreshing tabayashi's comment.
https://cloud.google.com/resource-manager/docs/organization-policy/restricting-service-accounts#disable_service_account_creation

upvoted 5 times

TNT87 Most Recent 1 year, 5 months ago

Selected Answer: D

Answer D
You can use the iam.disableServiceAccountCreation boolean constraint to disable the creation of new service accounts. This allows you to centrali
management of service accounts while not restricting the other permissions your developers have on projects.

upvoted 1 times

pskm12 1 year, 7 months ago

In the question, it is clearly mentioned that -> You want to accomplish this centrally across the organization. So, it would obviously be D

upvoted 1 times

gupta3 1 year, 8 months ago

Selected Answer: A

Are they not conflicting - restricting service account creation capability within production environments & enforcing policy across Org ?

upvoted 1 times

AzureDP900 1 year, 10 months ago

D is correct

upvoted 2 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: D

D. Use organization policy constraints/iam.disableServiceAccountCreation boolean to disable the creation of new service accounts.

upvoted 2 times

zellck 1 year, 11 months ago

Selected Answer: D

D is the answer.

upvoted 2 times

Question #96 *Topic 1*

You are the project owner for a regulated workload that runs in a project you own and manage as an Identity and Access Management (IAM) admin. For an upcoming audit, you need to provide access reviews evidence. Which tool should you use?

A. Policy Troubleshooter

B. Policy Analyzer

C. IAM Recommender

D. Policy Simulator

Correct Answer: *A*

Reference:

https://cloud.google.com/iam/docs/granting-changing-revoking-access

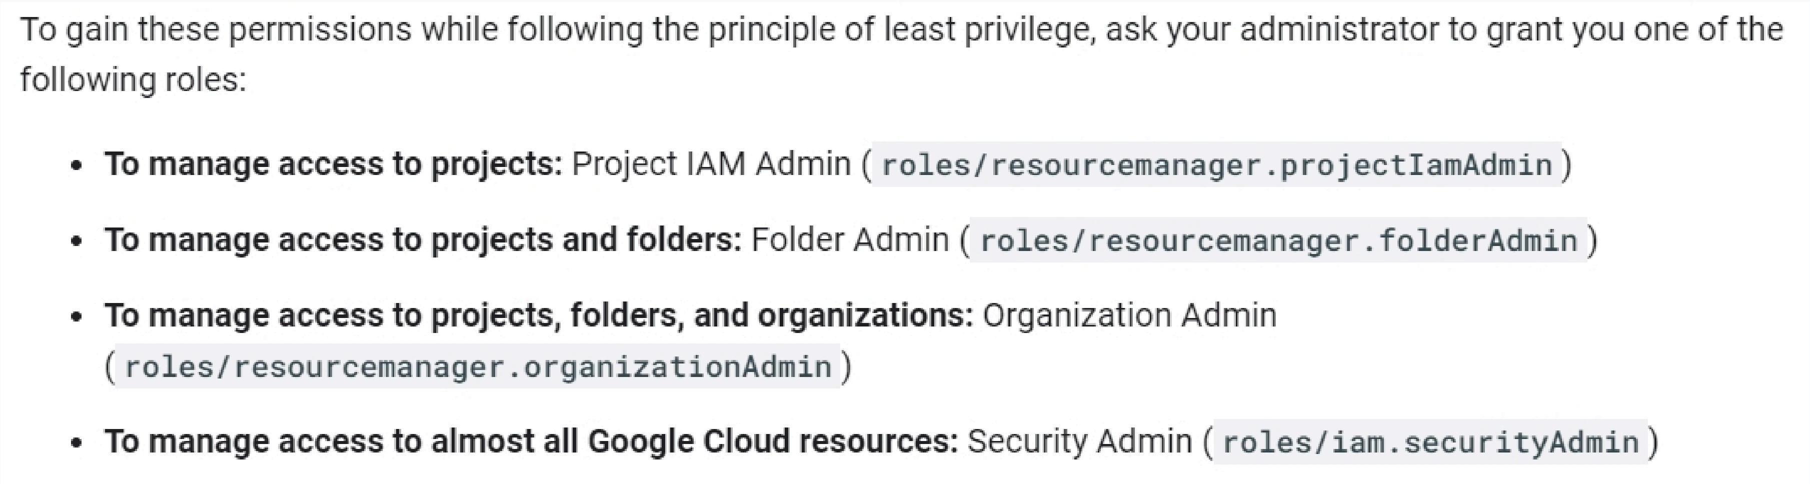

To gain these permissions while following the principle of least privilege, ask your administrator to grant you one of the following roles:

- **To manage access to projects:** Project IAM Admin (`roles/resourcemanager.projectIamAdmin`)
- **To manage access to projects and folders:** Folder Admin (`roles/resourcemanager.folderAdmin`)
- **To manage access to projects, folders, and organizations:** Organization Admin (`roles/resourcemanager.organizationAdmin`)
- **To manage access to almost all Google Cloud resources:** Security Admin (`roles/iam.securityAdmin`)

mouchu Highly Voted 2 years, 3 months ago

Answer = B

https://cloud.google.com/policy-intelligence/docs/policy-analyzer-overview

upvoted 10 times

sumundada Highly Voted 2 years, 1 month ago

Selected Answer: B

https://cloud.google.com/policy-intelligence/docs/policy-analyzer-overview

upvoted 5 times

rwintrob Most Recent 1 year, 6 months ago

B policy analyzer is the correct answer

upvoted 2 times

AzureDP900 1 year, 10 months ago

B policy analyzer is correct

upvoted 1 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: B

B. Policy Analyzer

upvoted 2 times

zellck 1 year, 11 months ago

Selected Answer: B

B is the answer.

https://cloud.google.com/policy-intelligence/docs/policy-analyzer-overview
Policy Analyzer lets you find out which principals (for example, users, service accounts, groups, and domains) have what access to which Google Cloud resources based on your IAM allow policies.

upvoted 3 times

cloudprincipal 2 years, 3 months ago

Selected Answer: B

https://cloud.google.com/policy-intelligence/docs/policy-analyzer-overview

upvoted 5 times

szl0144 2 years, 3 months ago

B is correct, guys

upvoted 4 times

Question #97 *Topic 1*

Your organization has implemented synchronization and SAML federation between Cloud Identity and Microsoft Active Directory. You want to reduce the risk of
Google Cloud user accounts being compromised. What should you do?

A. Create a Cloud Identity password policy with strong password settings, and configure 2-Step Verification with security keys in the Google Admin console.

B. Create a Cloud Identity password policy with strong password settings, and configure 2-Step Verification with verification codes via text or phone call in the Google Admin console.

C. Create an Active Directory domain password policy with strong password settings, and configure post-SSO (single sign-on) 2-Step Verification with security keys in the Google Admin console.

D. Create an Active Directory domain password policy with strong password settings, and configure post-SSO (single sign-on) 2-Step Verification with verification codes via text or phone call in the Google Admin console.

Correct Answer: *D*

Reference:
https://cloud.google.com/architecture/identity/federating-gcp-with-active-directory-introduction

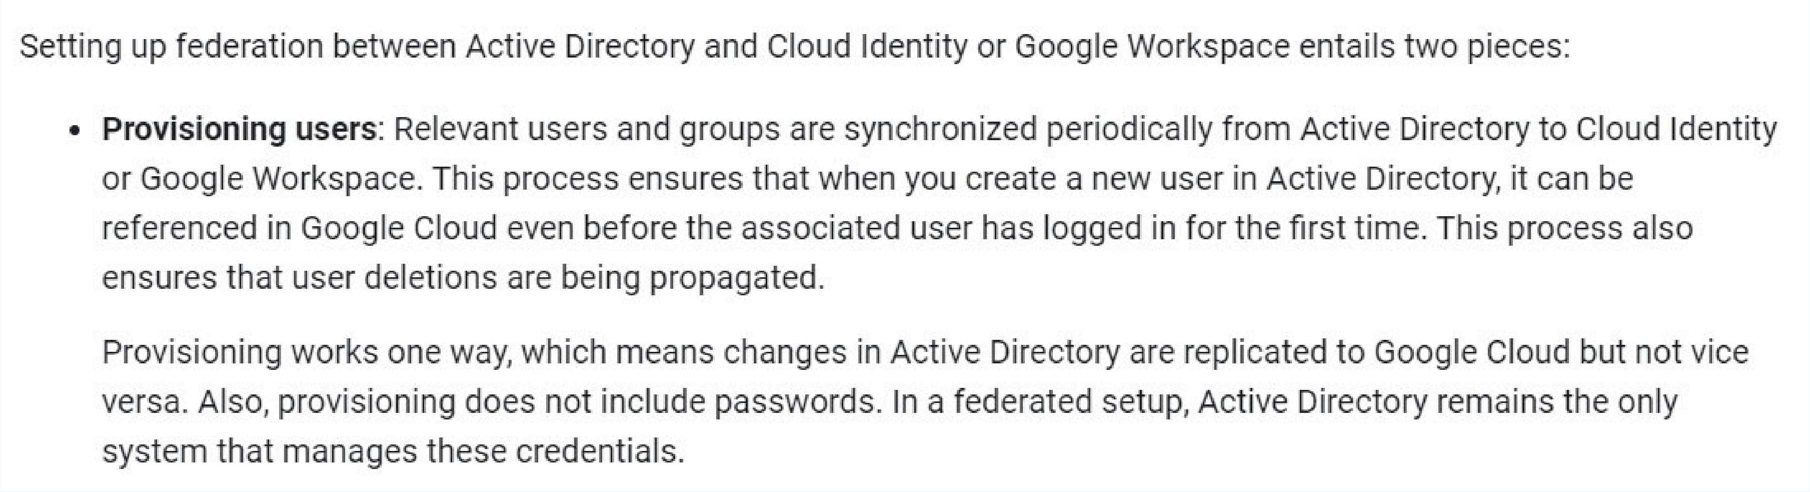

Setting up federation between Active Directory and Cloud Identity or Google Workspace entails two pieces:

- **Provisioning users**: Relevant users and groups are synchronized periodically from Active Directory to Cloud Identity or Google Workspace. This process ensures that when you create a new user in Active Directory, it can be referenced in Google Cloud even before the associated user has logged in for the first time. This process also ensures that user deletions are being propagated.

  Provisioning works one way, which means changes in Active Directory are replicated to Google Cloud but not vice versa. Also, provisioning does not include passwords. In a federated setup, Active Directory remains the only system that manages these credentials.

coco10k Highly Voted 1 year, 10 months ago

Answer C:
"We recommend against using text messages. The National Institute of Standards and Technology (NIST) no longer recommends SMS-based 2SV due to the hijacking risk from state-sponsored entities."

upvoted 5 times

gcpengineer 1 year, 3 months ago

user account doesnt need admin console access

upvoted 1 times

uiuiui Most Recent 10 months ago

Selected Answer: C

"C" Please

upvoted 2 times

[Removed] 1 year, 1 month ago

Selected Answer: C

"C"
Because it's federated access, the password policy stays with the origin IDP (Active Directory in this case) while the post-sso behavior/controls are in Google Cloud.
In terms of the actual second factor, security keys are far more secure than otp via text since those can be defeated through smishing or other types of attacks.

https://cloud.google.com/architecture/identity/federating-gcp-with-active-directory-introduction#implementing_federation
https://cloud.google.com/identity/solutions/enforce-mfa#use_security_keys

upvoted 4 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: C

C. Create an Active Directory domain password policy with strong password settings, and configure post-SSO (single sign-on) 2-Step Verification with security keys in the Google Admin console.

upvoted 3 times

jitu028 1 year, 11 months ago

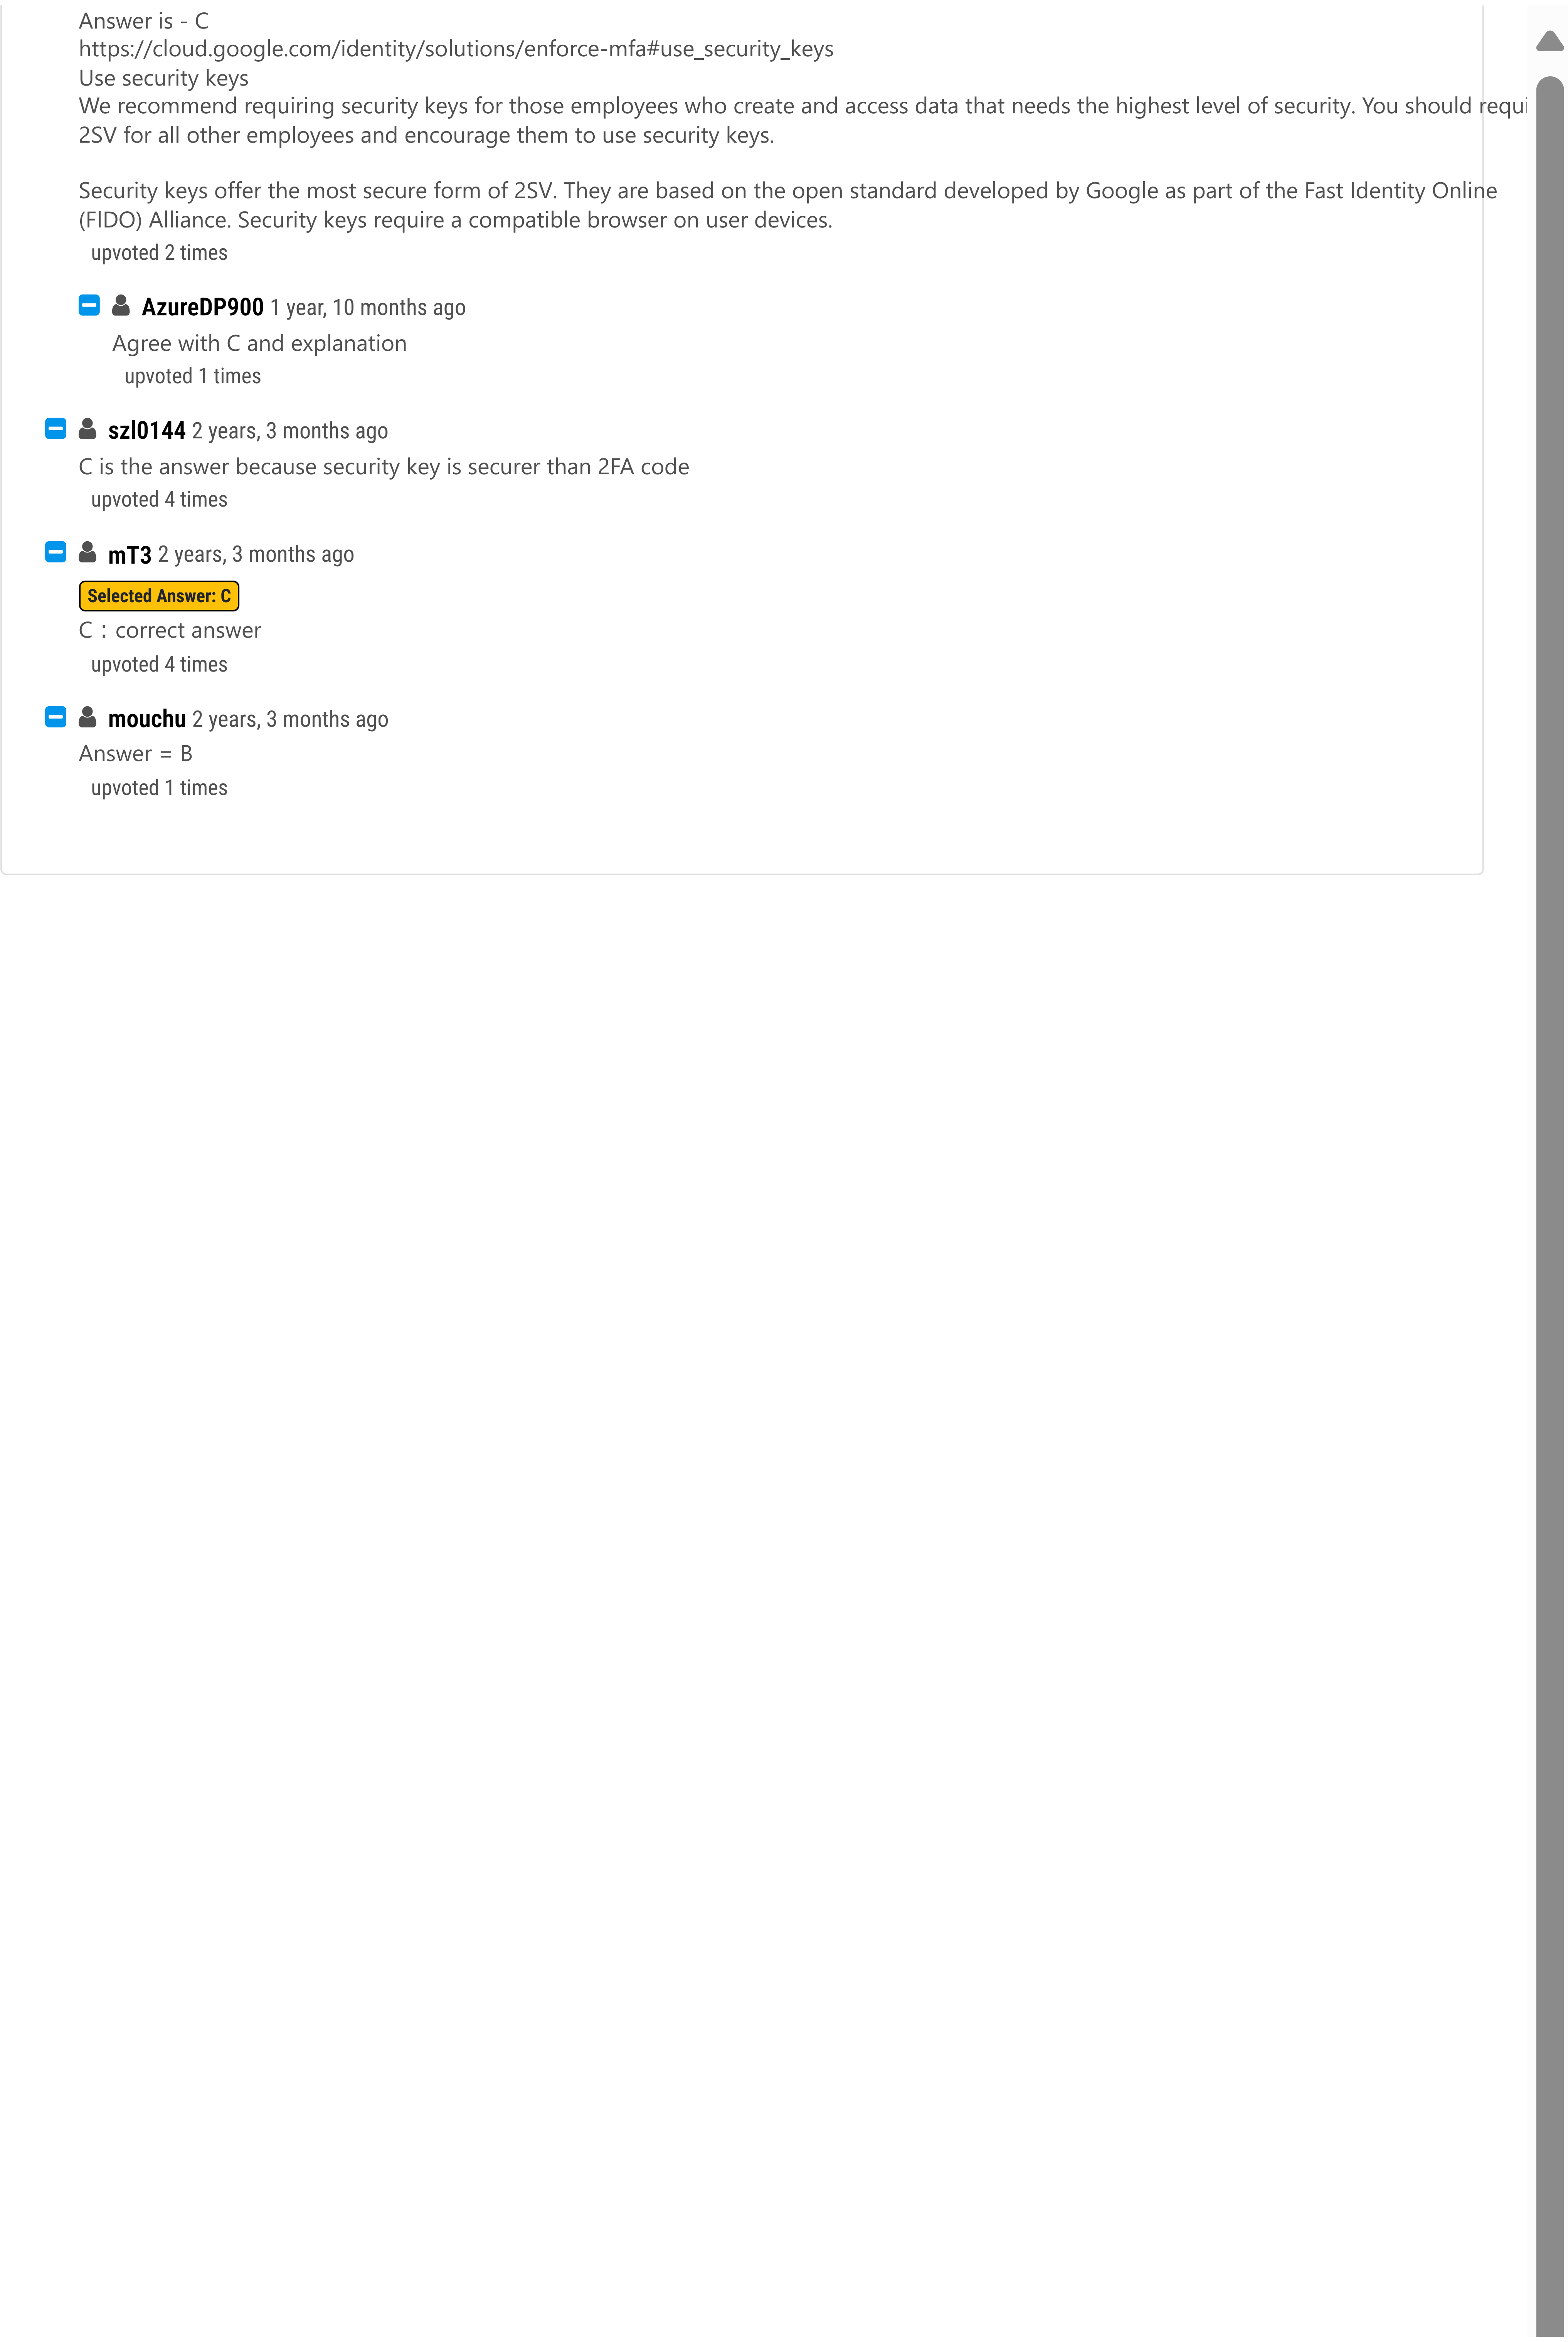

Answer is - C
https://cloud.google.com/identity/solutions/enforce-mfa#use_security_keys
Use security keys
We recommend requiring security keys for those employees who create and access data that needs the highest level of security. You should requi
2SV for all other employees and encourage them to use security keys.

Security keys offer the most secure form of 2SV. They are based on the open standard developed by Google as part of the Fast Identity Online (FIDO) Alliance. Security keys require a compatible browser on user devices.

upvoted 2 times

AzureDP900 1 year, 10 months ago

Agree with C and explanation

upvoted 1 times

szl0144 2 years, 3 months ago

C is the answer because security key is securer than 2FA code

upvoted 4 times

mT3 2 years, 3 months ago

Selected Answer: C

C：correct answer

upvoted 4 times

mouchu 2 years, 3 months ago

Answer = B

upvoted 1 times

Question #98 *Topic 1*

You have been tasked with implementing external web application protection against common web application attacks for a public application on Google Cloud.
You want to validate these policy changes before they are enforced. What service should you use?

A. Google Cloud Armor's preconfigured rules in preview mode

B. Prepopulated VPC firewall rules in monitor mode

C. The inherent protections of Google Front End (GFE)

D. Cloud Load Balancing firewall rules

E. VPC Service Controls in dry run mode

Correct Answer: *A*
Reference:
https://cloud.google.com/architecture/owasp-top-ten-mitigation

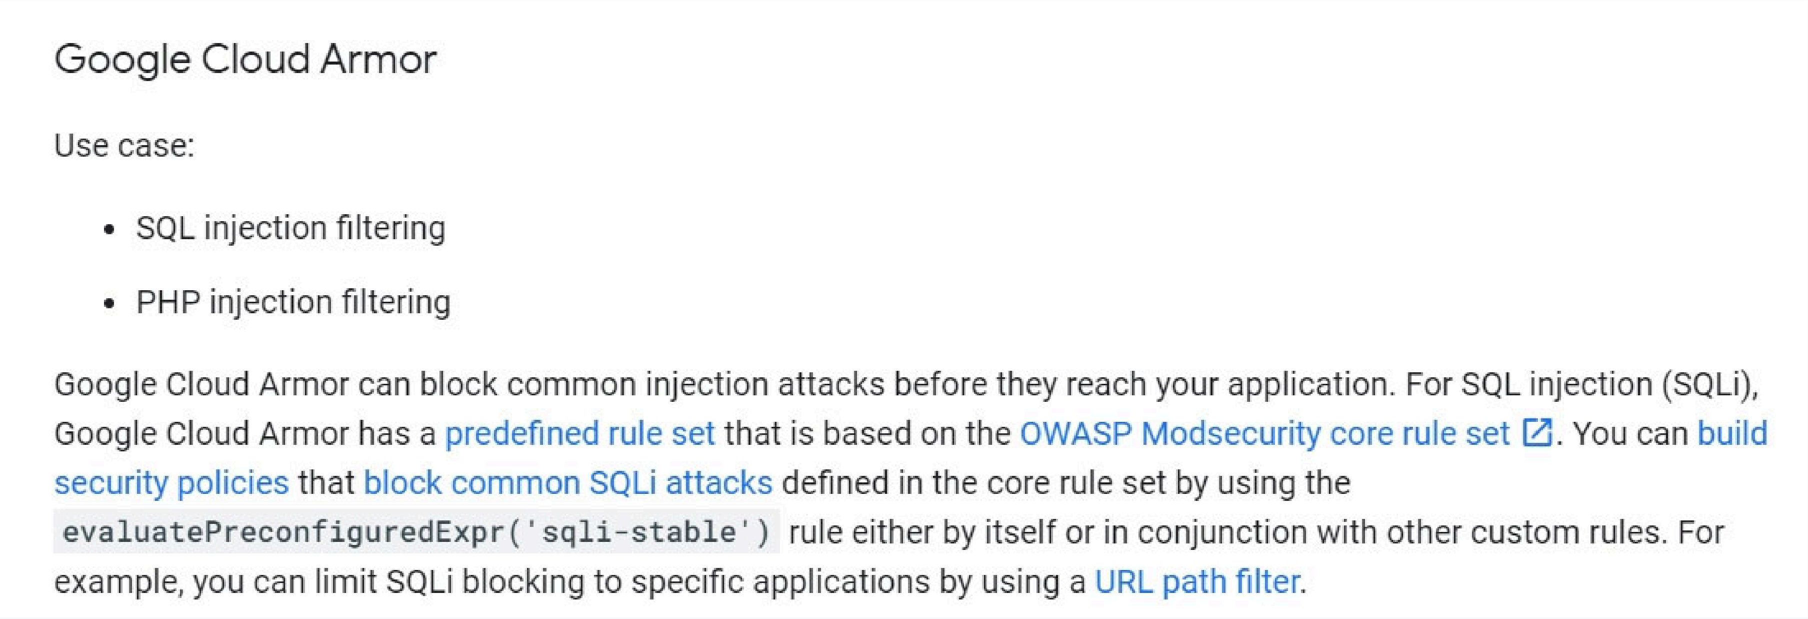

Google Cloud Armor

Use case:

- SQL injection filtering
- PHP injection filtering

Google Cloud Armor can block common injection attacks before they reach your application. For SQL injection (SQLi), Google Cloud Armor has a predefined rule set that is based on the OWASP Modsecurity core rule set ↗. You can build security policies that block common SQLi attacks defined in the core rule set by using the `evaluatePreconfiguredExpr('sqli-stable')` rule either by itself or in conjunction with other custom rules. For example, you can limit SQLi blocking to specific applications by using a URL path filter.

Tabayashi Highly Voted 2 years, 4 months ago
Answer is (A).

You can preview the effects of a rule without enforcing it. In preview mode, actions are noted in Cloud Monitoring. You can choose to preview individual rules in a security policy, or you can preview every rule in the policy.
https://cloud.google.com/armor/docs/security-policy-overview#preview_mode
upvoted 10 times

AzureDP900 1 year, 10 months ago
A is right
upvoted 1 times

tia_gll Most Recent 5 months, 2 weeks ago
Selected Answer: A
ans is A
upvoted 1 times

[Removed] 1 year, 1 month ago
Selected Answer: A
"A"
Web Application Firewall (Cloud Armor) is the answer here with preview mode.

https://cloud.google.com/armor/docs/security-policy-overview#preview_mode
upvoted 2 times

AwesomeGCP 1 year, 11 months ago
Selected Answer: A
A. Google Cloud Armor's preconfigured rules in preview mode
upvoted 2 times

sumundada 2 years, 1 month ago
Selected Answer: A
Answer is (A).
upvoted 2 times

Question #99 *Topic 1*

You are asked to recommend a solution to store and retrieve sensitive configuration data from an application that runs on Compute Engine. Which option should you recommend?

A. Cloud Key Management Service

B. Compute Engine guest attributes

C. Compute Engine custom metadata

D. Secret Manager

Correct Answer: *D*

Reference:

https://www.freecodecamp.org/news/google-cloud-platform-from-zero-to-hero/

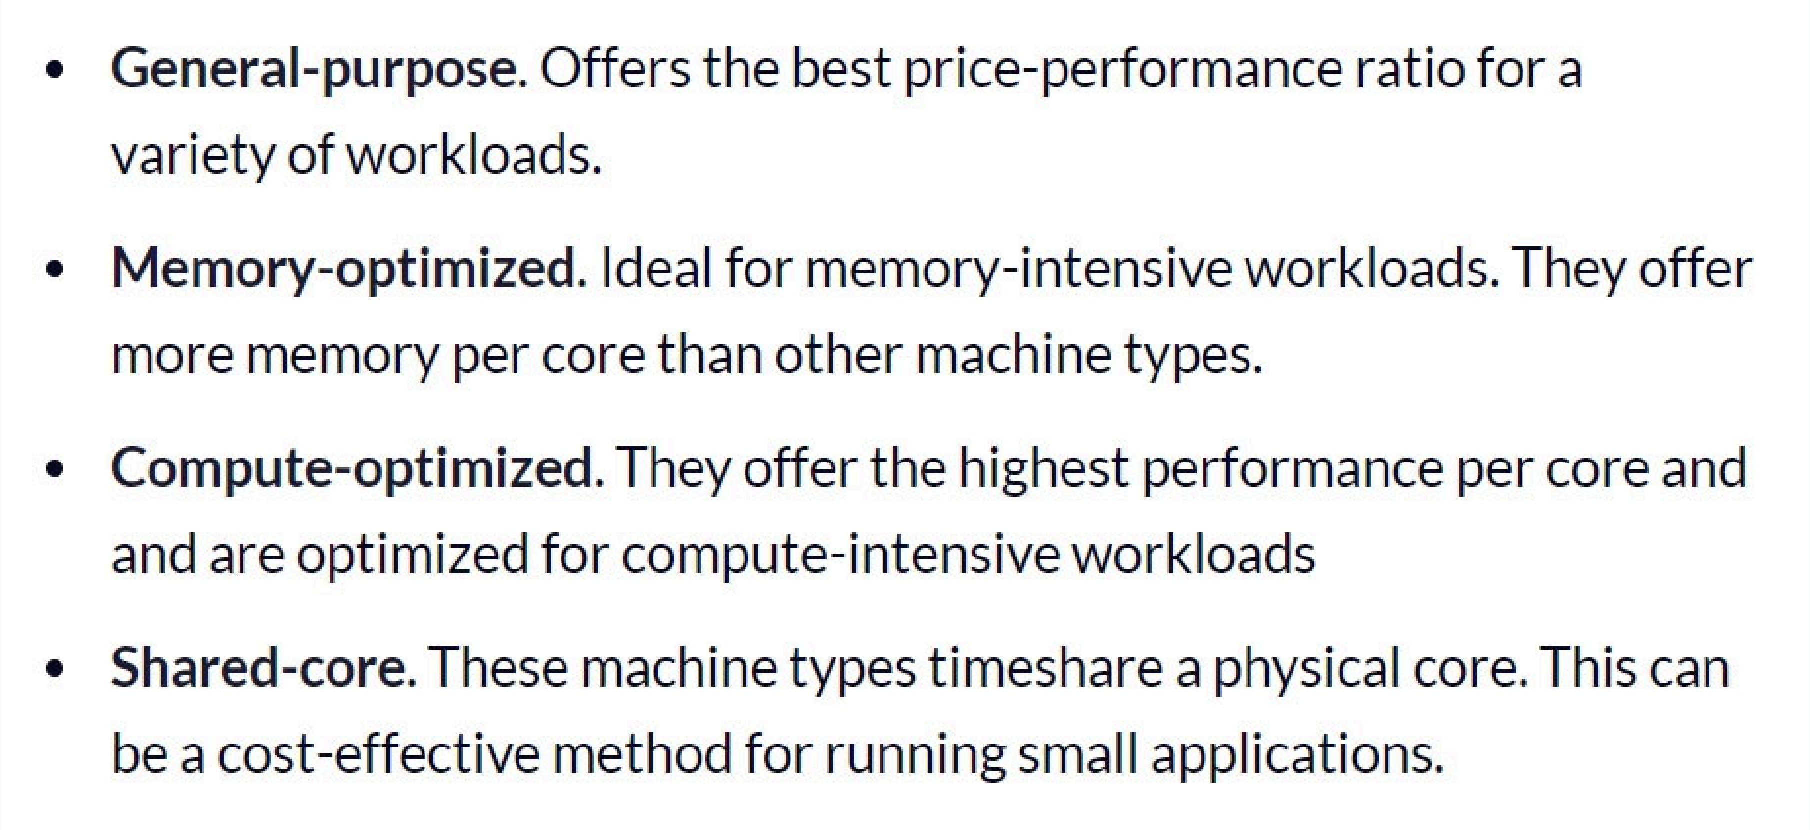

- **General-purpose.** Offers the best price-performance ratio for a variety of workloads.
- **Memory-optimized.** Ideal for memory-intensive workloads. They offer more memory per core than other machine types.
- **Compute-optimized.** They offer the highest performance per core and and are optimized for compute-intensive workloads
- **Shared-core.** These machine types timeshare a physical core. This can be a cost-effective method for running small applications.

Tabayashi Highly Voted 2 years, 4 months ago

Answer is (D).

Secret Manager is a secure and convenient storage system for API keys, passwords, certificates, and other sensitive data. Secret Manager provide
a central place and single source of truth to manage, access, and audit secrets across Google Cloud.
https://cloud.google.com/secret-manager

upvoted 13 times

cloudprincipal Highly Voted 2 years, 3 months ago

Selected Answer: D

You need a secrets management solution
https://cloud.google.com/secret-manager

upvoted 5 times

cloudprincipal 2 years, 3 months ago

Sorry, this should be C

upvoted 1 times

badrik 2 years, 3 months ago

sensitive information can never be stored/retrieved through custom meta data !

upvoted 4 times

tia_gll Most Recent 5 months, 2 weeks ago

Selected Answer: D

ans is D

upvoted 1 times

dija123 6 months ago

Selected Answer: D

Secret Manager

upvoted 1 times

[Removed] 1 year, 1 month ago

Selected Answer: D

"D"
There's ambiguity in the question in terms of what type of configuration data we're talking about and how large. Even though the compute metadata server can hold sensitive values like ssh keys, there are limitations with respect to how much data you can put in there (reference A below). Secret manager also has a size limit on how much you can store. (reference B below). However, secret manager is explicitly said to be a good use case for Sensitive Configuration information (reference C below) which makes it the preferred answer.

References:
A- https://cloud.google.com/compute/docs/metadata/setting-custom-metadata#limitations
B- https://cloud.google.com/secret-manager/quotas
C- https://cloud.google.com/secret-manager/docs/overview#secret_manager

upvoted 3 times

AzureDP900 1 year, 10 months ago

D is correct

upvoted 2 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: D

D. Secret Manager

upvoted 2 times

Question #100 *Topic 1*

You need to implement an encryption at-rest strategy that reduces key management complexity for non-sensitive data and protects sensitive data while providing the flexibility of controlling the key residency and rotation schedule. FIPS 140-2 L1 compliance is required for all data types. What should you do?

A. Encrypt non-sensitive data and sensitive data with Cloud External Key Manager.

B. Encrypt non-sensitive data and sensitive data with Cloud Key Management Service

C. Encrypt non-sensitive data with Google default encryption, and encrypt sensitive data with Cloud External Key Manager.

D. Encrypt non-sensitive data with Google default encryption, and encrypt sensitive data with Cloud Key Management Service.

Correct Answer: *B*

Chute5118 Highly Voted 2 years, 1 month ago

Selected Answer: D

Both B and D seem correct tbh. D might be "more correct" depending on the interpretation.

"reduces key management complexity for non-sensitive data" - Google default encryption
"protects sensitive data while providing the flexibility of controlling the key residency and rotation schedule" - Customer Managed Key

upvoted 6 times

AzureDP900 1 year, 10 months ago

I agree, D is right

upvoted 2 times

zellck Highly Voted 1 year, 11 months ago

Selected Answer: D

D is the answer.

upvoted 5 times

dija123 Most Recent 6 months ago

Selected Answer: D

D. Encrypt non-sensitive data with Google default encryption, and encrypt sensitive data with Cloud Key Management Service (KMS)

upvoted 1 times

MHD84 1 year ago

corrcet Answer is D, both KMS and default encryption are FIPS 140-2 L1 compliance https://cloud.google.com/kms/docs/key-management-service#choose

upvoted 3 times

[Removed] 1 year, 1 month ago

Selected Answer: D

"D"
Default encryption is Fips 140-2 L2 compliant (reference A below). Cloud KMS provides the rotation convenience desired (reference B below).

References:
A- https://cloud.google.com/docs/security/encryption/default-encryption
B- https://cloud.google.com/docs/security/key-management-deep-dive

upvoted 3 times

passex 1 year, 8 months ago

"reduces key management" & "FIPS 140-2 L1 compliance is required for all data types" - strongly suggests answer B

upvoted 1 times

rrvv 1 year, 11 months ago

As FIPS 140-2 L1 compliance is required for all types of data, Cloud KMS should be used to manage encryption. Correct answer is B

https://cloud.google.com/docs/security/key-management-deep-dive#software-protection-level:~:text=The%20Cloud%20KMS%20binary%20is%20built%20against%20FIPS%20140%2D2%20Level%201%E2%80%93validated%20Cryptogra hic%20Primitives%20of%20this%20module

upvoted 1 times

sumundada 2 years, 1 month ago

Selected Answer: D

Google uses a common cryptographic library, Tink, which incorporates our FIPS 140-2 Level 1 validated module, BoringCrypto, to implement encryption consistently across almost all Google Cloud products. To provideflexibility of controlling the key residency and rotation schedule, use google provided key for non-sensitive and encrypt sensitive data with Cloud Key Management Service

upvoted 3 times

nacying 2 years, 2 months ago

Selected Answer: B

base on "FIPS 140-2 L1 compliance is required for all data types"

upvoted 3 times

cloudprincipal 2 years, 3 months ago

Selected Answer: D

KMS is ok for fips 140-2 level 1
https://cloud.google.com/docs/security/key-management-deep-dive#platform-overview

upvoted 2 times

> **cloudprincipal** 2 years, 2 months ago
>
> Regarding FIPS 140-2 level 1 and GCP default encryption:
>
> Google Cloud uses a FIPS 140-2 validated Level 1 encryption module (certificate 3318) in our production environment.
>
> https://cloud.google.com/docs/security/encryption/default-encryption?hl=en#encryption_of_data_at_rest
>
> upvoted 2 times

mikesp 2 years, 3 months ago

In my opinion, the answer is B. The question says that it is necessary to control "key residency and rotation schedule" for both types of data. Default encryption at rest does not provide that but Cloud KMS does. Furthermore, Cloud KMS is FIPS140-2 level 1.
https://cloud.google.com/docs/security/key-management-deep-dive

upvoted 3 times

> **csrazdan** 1 year, 8 months ago
>
> The answer is D.
> 1. reduce key management complexity for non-sensitive data --> Google Managed key
> 2. protects sensitive data while providing the flexibility of controlling the key residency and rotation schedule --> KMS
>
> upvoted 1 times

szl0144 2 years, 3 months ago

D is the wander

upvoted 3 times

mouchu 2 years, 3 months ago

Answer = D

upvoted 3 times

Question #101 Topic 1

Your company wants to determine what products they can build to help customers improve their credit scores depending on their age range. To achieve this, you need to join user information in the company's banking app with customers' credit score data received from a third party. While using this raw data will allow you to complete this task, it exposes sensitive data, which could be propagated into new systems.
This risk needs to be addressed using de-identification and tokenization with Cloud Data Loss Prevention while maintaining the referential integrity across the database. Which cryptographic token format should you use to meet these requirements?

A. Deterministic encryption

B. Secure, key-based hashes

C. Format-preserving encryption

D. Cryptographic hashing

Correct Answer: *B*
Reference:
https://cloud.google.com/blog/products/identity-security/take-charge-of-your-data-how-tokenization-makes-data-usable-without-sacrificing-privacy

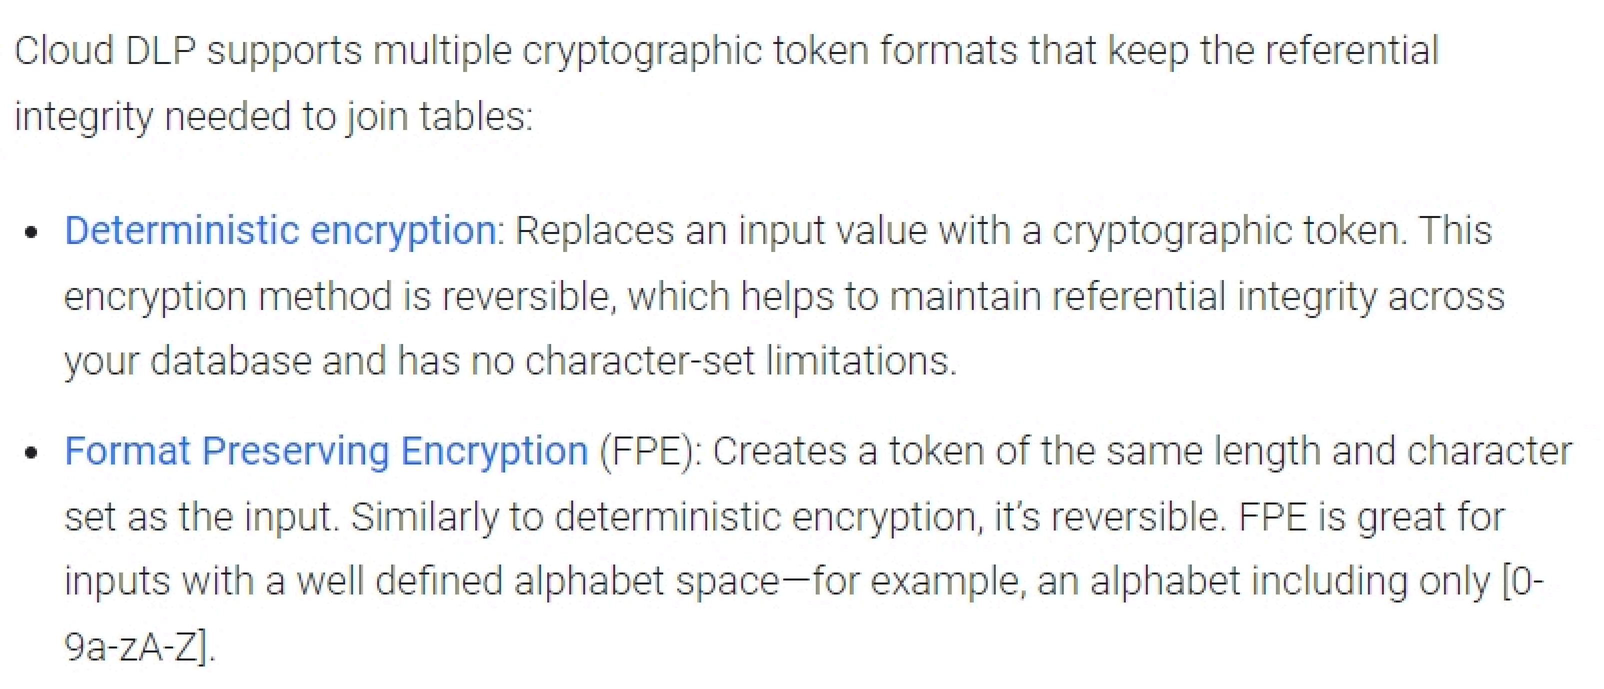

Cloud DLP supports multiple cryptographic token formats that keep the referential integrity needed to join tables:

- Deterministic encryption: Replaces an input value with a cryptographic token. This encryption method is reversible, which helps to maintain referential integrity across your database and has no character-set limitations.
- Format Preserving Encryption (FPE): Creates a token of the same length and character set as the input. Similarly to deterministic encryption, it's reversible. FPE is great for inputs with a well defined alphabet space—for example, an alphabet including only [0-9a-zA-Z].

mT3 Highly Voted 2 years, 3 months ago
Selected Answer: A
"This encryption method is reversible, which helps to maintain referential integrity across your database and has no character-set limitations."
https://cloud.google.com/blog/products/identity-security/take-charge-of-your-data-how-tokenization-makes-data-usable-without-sacrificing-privacy
upvoted 11 times

[Removed] 1 year, 1 month ago
I meant both A and C not A and D.
upvoted 1 times

AzureDP900 1 year, 10 months ago
A is right
upvoted 1 times

rsamant Most Recent 9 months ago
D Cryptogrpahic hashing as it maintains refenrtial integrity and not reversible https://cloud.google.com/dlp/docs/pseudonymization
upvoted 2 times

Xoxoo 11 months, 2 weeks ago
Selected Answer: A
To meet the requirements of de-identifying and tokenizing sensitive data while maintaining referential integrity across the database, you should use "Deterministic encryption."

Deterministic encryption is a form of encryption where the same input value consistently produces the same encrypted output (token). This ensur
referential integrity because the same original value will always result in the same token, allowing you to link and join data across different system
or databases while still protecting sensitive information.

Format-preserving encryption is a specific form of deterministic encryption that preserves the format and length of the original data, which can b
useful for maintaining data structures and relationships.

So, the correct option is:

A. Deterministic encryption

upvoted 2 times

[Removed] 1 year, 1 month ago

Selected Answer: A

"A"
Requirements are reversible while maintaining referential integrity. Both A and D meet this requirement however D has input limitations. Therefor
A is a better answer.

https://cloud.google.com/dlp/docs/transformations-reference#transformation_methods

upvoted 1 times

danidee111 1 year, 2 months ago

This is a poor question and not enough data is provided to determine which Tokenization method should be selected. There are three methods fo
Tokenization (also referred to as Pseudonymization). See: https://cloud.google.com/dlp/docs/transformations-reference#crypto and each methoc
maintains referential integrity See: https://www.youtube.com/watch?v=h0BnA7R8vg4. Thus, you'd need to know whether it needs to be reversible
format preserving to confidentially select an answer..

upvoted 3 times

gcpengineer 1 year, 3 months ago

Selected Answer: A

https://cloud.google.com/blog/products/identity-security/take-charge-of-your-data-how-tokenization-makes-data-usable-without-sacrificing-privacy

upvoted 1 times

passex 1 year, 8 months ago

"Deterministic encryption" is too wide definition, the key phrase is "Which cryptographic token format " so th answer is "Format-preserving encryption" - where Referential integrity is assured (...allows for records to maintain their relationship ....ensures that connections between values (and, with structured data, records) are preserved, even across tables)

upvoted 1 times

gcpengineer 1 year, 3 months ago

A is the ans. https://cloud.google.com/blog/products/identity-security/take-charge-of-your-data-how-tokenization-makes-data-usable-without-sacrificing-privacy

upvoted 1 times

PST21 1 year, 8 months ago

Cryptographic uses strings , it asks to use tokenization and hence deterministic is better than FPE hence A

upvoted 1 times

gcpengineer 1 year, 3 months ago

both create tokens, the FPE is more used where u have format [0-9a-za-Z]

upvoted 1 times

Littleivy 1 year, 9 months ago

Selected Answer: D

Though it's not clear, but, to prevent from data leak, it's better to have a non-reversible method as analysts don't need re-identification

upvoted 1 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: A

A. Deterministic encryption

upvoted 1 times

zellck 1 year, 11 months ago

Selected Answer: A

A is the answer.

https://cloud.google.com/dlp/docs/pseudonymization
FPE provides fewer security guarantees compared to other deterministic encryption methods such as AES-SIV.
For these reasons, Google strongly recommends using deterministic encryption with AES-SIV instead of FPE for all security sensitive use cases.

Other methods like deterministic encryption using AES-SIV provide these stronger security guarantees and are recommended for tokenization us
cases unless length and character set preservation are strict requirements—for example, for backward compatibility with a legacy data system.

upvoted 4 times

piyush_1982 2 years, 1 month ago

Selected Answer: A

This question is taken from the exact scenario described in this link
https://cloud.google.com/blog/products/identity-security/take-charge-of-your-data-how-tokenization-makes-data-usable-without-sacrificing-privacy

upvoted 1 times

Chute5118 2 years, 1 month ago

Selected Answer: D

Both "Deterministic" and "format preserving" are key-based hashes (and reversible).
It's not clear from the question, but doesn't look like we need it to be reversible.
All of them maintain referential integrity
https://cloud.google.com/architecture/de-identification-re-identification-pii-using-cloud-dlp#method_selection

upvoted 1 times

cloudprincipal 2 years, 3 months ago

Selected Answer: D

preserve referential integrity and ensure that no re-identification is possible

https://cloud.google.com/dlp/docs/pseudonymization#supported-methods

upvoted 1 times

> **cloudprincipal** 2 years, 2 months ago
>
> forget it, it should be A.
>
> upvoted 1 times

Taliesyn 2 years, 3 months ago

Selected Answer: D

Cryptographic hash (CryptoHashConfig) maintains referential integrity.
"Determinist encryption" is not a transformation method.
https://cloud.google.com/dlp/docs/transformations-reference

upvoted 2 times

gionny 2 years, 4 months ago

A correct - https://cloud.google.com/blog/products/identity-security/take-charge-of-your-data-how-tokenization-makes-datausable-without-sacrificing-privacy

upvoted 3 times

Question #102 *Topic 1*

An office manager at your small startup company is responsible for matching payments to invoices and creating billing alerts. For compliance reasons, the office manager is only permitted to have the Identity and Access Management (IAM) permissions necessary for these tasks. Which two IAM roles should the office manager have? (Choose two.)

A. Organization Administrator

B. Project Creator

C. Billing Account Viewer

D. Billing Account Costs Manager

E. Billing Account User

Correct Answer: *AE*

Reference:

https://cloud.google.com/iam/docs/job-functions/billing

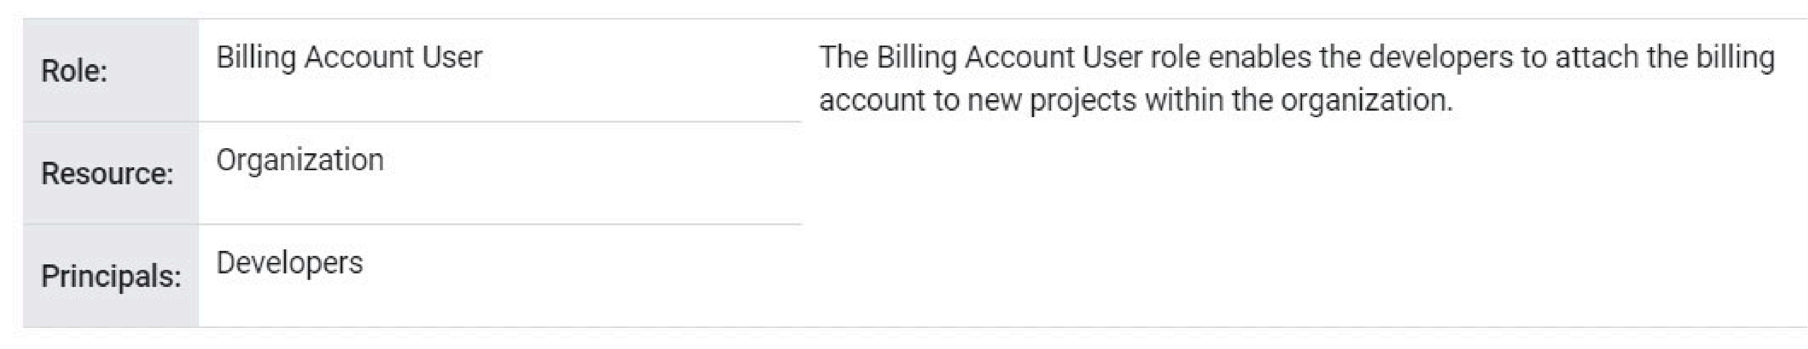

Role:	Billing Account User	The Billing Account User role enables the developers to attach the billing account to new projects within the organization.
Resource:	Organization	
Principals:	Developers	

mT3 Highly Voted 👍 2 years, 3 months ago

Selected Answer: CD

Ans C,D.
C. Billing Account Viewer :responsible for matching payments to invoices
https://cloud.google.com/billing/docs/how-to/get-invoice#required-permissions
Access billing documents:"Billing Account Administrator" or "Billing Account Viewer"
D. Billing Account Costs Manager : creating billing alerts
https://cloud.google.com/billing/docs/how-to/budgets-notification-recipients
"To create or modify a budget for your Cloud Billing account, you need the Billing Account Costs Manager role or the Billing Account Administrat role on the Cloud Billing account."
and "If you want the recipients of the alert emails to be able to view the budget, email recipients need permissions on the Cloud Billing account. / a minimum, ensure email recipients are added to the Billing Account Viewer role on the Cloud Billing account that owns the budget. See View a li of budgets for additional information."

upvoted 15 times

AzureDP900 1 year, 10 months ago

CD is right

upvoted 3 times

GHOST1985 1 year, 9 months ago

the link you post talking about Permissions required to ACCESS billing documentsn not to link project to a billing account you should have the Billing Account User role, the good answer is D,E

upvoted 1 times

Taliesyn Highly Voted 👍 2 years, 3 months ago

Selected Answer: CD

Billing Account Costs Administrator to create budgets (aka alerts)
Billing Account Viewer to view costs (to be able to match them to invoices)

upvoted 6 times

rottzy Most Recent ⏱ 11 months, 2 weeks ago

Billing Account Costs Manager - does not exist! ?!

upvoted 1 times

winston9 6 months, 3 weeks ago

yes, it does: https://cloud.google.com/iam/docs/understanding-roles#billing.costsManager

upvoted 1 times

desertlotus1211 1 year ago

Answer: CD
https://cloud.google.com/billing/docs/how-to/budgets

upvoted 1 times

[Removed] 1 year, 1 month ago

Selected Answer: CD

C,D
BA Viewer to see spend info and BA Costs Manager to manage costs, create budgets and alerts
BA User and BA Admin have permissions related to linking projects to billing etc. which are not needed.
https://cloud.google.com/billing/docs/how-to/billing-access#ba-viewer
https://cloud.google.com/billing/docs/how-to/billing-access

upvoted 2 times

GHOST1985 1 year, 10 months ago

Selected Answer: DE

Billing Account User: This role has very restricted permissions, so you can grant it broadly. When granted in combination with Project Creator, the two roles allow a user to create new projects linked to the billing account on which the Billing Account User role is granted. Or, when granted in combination with the Project Billing Manager role, the two roles allow a user to link and unlink projects on the billing account on which the Billinc Account User role is granted.

Billing Account Costs Manager: Create, edit, and delete budgets, view billing account cost information and transactions, and manage the export c billing cost data to BigQuery. Does not confer the right to export pricing data or view custom pricing in the Pricing page. Also, does not allow the linking or unlinking of projects or otherwise managing the properties of the billing account

upvoted 3 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: CD

C. Billing Account Viewer
D. Billing Account Costs Manager

upvoted 2 times

zellck 1 year, 11 months ago

Selected Answer: CD

CD is the answer.

https://cloud.google.com/billing/docs/how-to/billing-access#overview-of-cloud-billing-roles-in-cloud-iam

Billing Account Costs Manager (roles/billing.costsManager)
- Manage budgets and view and export cost information of billing accounts (but not pricing information)

Billing Account Viewer (roles/billing.viewer)
- View billing account cost information and transactions.

upvoted 3 times

Question #103 *Topic 1*

You are designing a new governance model for your organization's secrets that are stored in Secret Manager. Currently, secrets for Production and Non-

Production applications are stored and accessed using service accounts. Your proposed solution must:

➾ Provide granular access to secrets

➾ Give you control over the rotation schedules for the encryption keys that wrap your secrets

➾ Maintain environment separation

➾ Provide ease of management

Which approach should you take?

A. 1. Use separate Google Cloud projects to store Production and Non-Production secrets. 2. Enforce access control to secrets using project-level identity and Access Management (IAM) bindings. 3. Use customer-managed encryption keys to encrypt secrets.

B. 1. Use a single Google Cloud project to store both Production and Non-Production secrets. 2. Enforce access control to secrets using secret-level Identity and Access Management (IAM) bindings. 3. Use Google-managed encryption keys to encrypt secrets.

C. 1. Use separate Google Cloud projects to store Production and Non-Production secrets. 2. Enforce access control to secrets using secret-level Identity and Access Management (IAM) bindings. 3. Use Google-managed encryption keys to encrypt secrets.

D. 1. Use a single Google Cloud project to store both Production and Non-Production secrets. 2. Enforce access control to secrets using project-level Identity and Access Management (IAM) bindings. 3. Use customer-managed encryption keys to encrypt secrets.

Correct Answer: *A*

mT3 Highly Voted 2 years, 3 months ago

Selected Answer: A

Correct. Ans A.
Provide granular access to secrets: 2.Enforce access control to secrets using project-level identity and Access Management (IAM) bindings.
Give you control over the rotation schedules for the encryption keys that wrap your secrets: 3. Use customer-managed encryption keys to encrypt secrets.
Maintain environment separation: 1. Use separate Google Cloud projects to store Production and Non-Production secrets.

upvoted 13 times

mikesp 2 years, 3 months ago

It is possible to grant IAM bindind to secret-level which is more granular than project-level but considering that it is necessary to manage encryption keys life-cycle, then the answer is A due to C does not allow that.

upvoted 4 times

AzureDP900 1 year, 10 months ago

Yes , A is right

upvoted 1 times

Medofree Highly Voted 2 years, 3 months ago

None of the answers are correct, here is why :

➾ Provide granular access to secrets => 2. Enforce access control to secrets using secret-level (and not project-level)
➾ Give you control over the rotation schedules for the encryption keys that wrap your secrets => 3. Use customer-managed encryption keys to encrypt secrets.
➾ Maintain environment separation => 1. Use separate Google Cloud projects to store Production and Non-Production secrets
➾ Provide ease of management => 3. Use Google-managed encryption keys to encrypt secrets. (could be in contradiction with Give you control over the rotation schedules....)

It should be an E answer :

E. 1. Use separate Google Cloud projects to store Production and Non-Production secrets. 2. Enforce access control to secrets using secret-level identity and Access Management (IAM) bindings. 3. Use customer-managed encryption keys to encrypt secrets.

upvoted 5 times

desertlotus1211 1 year ago

That's Answer A....

upvoted 1 times

glb2 Most Recent 5 months, 3 weeks ago

Selected Answer: C

I think C is correct.
Secrets granular management, separate projects and keys managements into google.

upvoted 1 times

[Removed] 8 months, 2 weeks ago

Selected Answer: C

For me this is answer C.
It provides granular access control at the secret level. Option A provides project-level IAM bindings and not secret level.
While it uses Google-managed keys (offering less control over rotation), it simplifies management and still maintains a good security posture.
It maintains environment separation by using different projects for Production and Non-Production.
Balances between ease of management and security, though slightly more complex due to separate projects.

upvoted 2 times

glb2 5 months, 3 weeks ago

I think the same.

upvoted 1 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: A

A. 1. Use separate Google Cloud projects to store Production and Non-Production secrets. 2. Enforce access control to secrets using project-level identity and Access Management (IAM) bindings. 3. Use customer-managed encryption keys to encrypt secrets.

upvoted 3 times

Question #104 *Topic 1*

You are a security engineer at a finance company. Your organization plans to store data on Google Cloud, but your leadership team is worried about the security of their highly sensitive data. Specifically, your company is concerned about internal Google employees' ability to access your company's data on Google Cloud.
What solution should you propose?

A. Use customer-managed encryption keys.

B. Use Google's Identity and Access Management (IAM) service to manage access controls on Google Cloud.

C. Enable Admin activity logs to monitor access to resources.

D. Enable Access Transparency logs with Access Approval requests for Google employees.

Correct Answer: *D*

zellck Highly Voted 1 year, 11 months ago

Selected Answer: D

D is the answer

upvoted 5 times

Sammydp202020 Highly Voted 1 year, 6 months ago

Selected Answer: D

D

https://cloud.google.com/access-transparency
https://cloud.google.com/cloud-provider-access-management/access-transparency/docs/overview

upvoted 5 times

Xoxoo Most Recent 11 months, 3 weeks ago

Selected Answer: D

To address your organization's concerns about the security of highly sensitive data stored on Google Cloud, you can propose the following solution:

D. Enable Access Transparency logs with Access Approval requests for Google employees. This solution provides an additional layer of control and visibility over your cloud provider by enabling you to monitor and audit the actions taken by Google personnel when accessing your content. Access Transparency logs capture the actions performed by Google Cloud administrators, allowing you to maintain an audit trail and verify cloud provider access. Access Approval requests allow you to approve or dismiss requests for access by Google employees working to support your service. By combining these features, you can gain greater oversight and control over your sensitive data on Google Cloud.

Please note that this is a high-level recommendation, and it is important to evaluate your specific requirements and consult the official Google Cloud documentation for detailed implementation guidance.

upvoted 3 times

passex 1 year, 8 months ago

Answer D - but, for "highly sensitive data" CMEK seems to be reasonable option but much easiest way is to use Transparency Logs

upvoted 1 times

PATILDXB 1 year, 8 months ago

B is the correct answer. IAM Privileges provide fine grain controls based on the users function

upvoted 1 times

Littleivy 1 year, 9 months ago

Selected Answer: A

Use customer-managed key to encrypt data by yourself

upvoted 2 times

Littleivy 1 year, 9 months ago

D should be the answer on second thought

upvoted 2 times

AzureDP900 1 year, 10 months ago

D is correct

upvoted 3 times

jitu028 1 year, 11 months ago

Answer is D

https://cloud.google.com/access-transparency

Access approval
Explicitly approve access to your data or configurations on Google Cloud. Access Approval requests, when combined with Access Transparency logs, can be used to audit an end-to-end chain from support ticket to access request to approval, to eventual access.

upvoted 4 times

Question #105 *Topic 1*

You want to use the gcloud command-line tool to authenticate using a third-party single sign-on (SSO) SAML identity provider. Which options are necessary to ensure that authentication is supported by the third-party identity provider (IdP)? (Choose two.)

A. SSO SAML as a third-party IdP

B. Identity Platform

C. OpenID Connect

D. Identity-Aware Proxy

E. Cloud Identity

Correct Answer: *AC*
Reference:
https://cloud.google.com/identity/solutions/enable-sso

ExamQnA Highly Voted 2 years, 3 months ago

Selected Answer: AE

Third-party identity providers
If you have a third-party IdP, you can still configure SSO for third-party apps in the Cloud Identity catalog. User authentication occurs in the third-party IdP, and Cloud Identity manages the cloud apps.

To use Cloud Identity for SSO, your users need Cloud Identity accounts. They sign in through your third-party IdP or using a password on their Cloud Identity accounts.
https://cloud.google.com/identity/solutions/enable-sso

upvoted 22 times

AzureDP900 1 year, 10 months ago

A, E is right

upvoted 5 times

piyush_1982 Highly Voted 2 years, 1 month ago

Selected Answer: AC

I think the correct answer is A and C.

The questions asks about what is required with third-party IdP to authenticate the gcloud commands.
So the gcloud command requests goes to GCP. Since GCP is integrated with Third-party IdP for authentication gcloud command needs to be authenticated with third-party IdP.

This can be achieved if ThridPaty IdP supports SAML and OIDC protocols .

upvoted 16 times

Mr_MIXER007 Most Recent 3 days, 18 hours ago

Selected Answer: AE

Selected Answer: AE

upvoted 1 times

3d9563b 1 month, 1 week ago

Selected Answer: AE

SSO SAML as a third-party IdP: This option ensures that the authentication mechanism used is SAML, which is required for third-party IdP integration.
Cloud Identity: This provides the underlying infrastructure to integrate and manage identities with third-party SAML IdPs, enabling SSO authentication.

upvoted 1 times

dija123 6 months ago

Selected Answer: CE

C. OpenID Connect
E. Cloud Identity
A. SSO SAML as a third-party IdP: While it accurately describes the desired authentication but It represents the outcome we want to achieve, not the solution itself.

upvoted 2 times

oezgan 5 months, 2 weeks ago

Gemini says: While SAML is a common protocol for SSO, it's not directly usable by gcloud for authentication. So it cant be A.

upvoted 2 times

mjcts 6 months, 4 weeks ago

Selected Answer: AE

OpenID is a different SSO protocol. We need SAML.

upvoted 2 times

Andras2k 8 months ago

Selected Answer: AE

It specifically requires the SAML protocol. OpenID is another SSO protocol.

upvoted 2 times

ymkk 1 year ago

Selected Answer: AE

Options B, C, and D are not directly related to setting up authentication using a third-party SSO SAML identity provider. Identity Platform (option B) is a service for authentication and user management, OpenID Connect (option C) is another authentication protocol, and Identity-Aware Proxy (option D) is a service for managing access to Google Cloud resources but is not specifically related to SSO SAML authentication with a third-part IdP.

upvoted 2 times

pfilourenco 1 year, 1 month ago

Selected Answer: AE

AE is the correct

upvoted 2 times

[Removed] 1 year, 1 month ago

Selected Answer: AE

"A,E"
The requirement is for an SSO - SAML solution with a third party IDP.
A- This is correct because it provides the right type of 3rd party partners.
B - Not sufficient because not any IDP will suffice. Must be able to support SAML and SSO.
C- OIDC is an option by not critical or a hard requirement. The questions asks about what is "..necessary..".
D- IAP is not related to authentication mechanism but rather authorization. This is not the use case for it.
E- This is needed on the receiving end in GCP to collaborate with 3rd party IDP (that has SAML SSO)

https://cloud.google.com/identity/solutions/enable-sso

upvoted 1 times

keymson 1 year, 4 months ago

OpenID Connect has to be there. so A and C

upvoted 1 times

testgcptestgcp 1 year, 3 months ago

Cloud Identity does not have to be there? Why?

upvoted 2 times

alleinallein 1 year, 5 months ago

Selected Answer: AC

Open ID seems to be necessary

upvoted 3 times

bruh_1 1 year, 5 months ago

A. SSO SAML as a third-party IdP: This option is necessary because it specifies that you want to use SAML-based SSO with a third-party IdP.

C. OpenID Connect: This option is necessary to ensure that the third-party IdP supports OpenID Connect, which is a protocol for authentication a authorization.

Therefore, the correct options are A and C.

upvoted 3 times

TNT87 1 year, 5 months ago

Selected Answer: AC

https://cloud.google.com/certificate-authority-service/docs/tutorials/using-3pi-with-reflection#set-up-wip

https://cloud.google.com/identity/solutions/enable-sso#solutions

Nothing supports E to satisfy the requirements othe question

upvoted 2 times

Sammydp202020 1 year, 6 months ago

Selected Answer: AE

AE

https://cloud.google.com/identity/solutions/enable-sso
Third-party identity providers
If you have a third-party IdP, you can still configure SSO for third-party apps in the Cloud Identity catalog. User authentication occurs in the third-

party IdP, and Cloud Identity manages the cloud apps.

To use Cloud Identity for SSO, your users need Cloud Identity accounts. They sign in through your third-party IdP or using a password on their Cloud Identity accounts.

upvoted 2 times

Littleivy 1 year, 9 months ago

Selected Answer: AC

answer is A and C.

upvoted 2 times

Littleivy 1 year, 9 months ago

Selected Answer: AC

To provide users with SSO-based access to selected cloud apps, Cloud Identity as your IdP supports the OpenID Connect (OIDC) and Security Assertion Markup Language 2.0 (SAML) protocols.

https://cloud.google.com/identity/solutions/enable-sso

upvoted 4 times

gcpengineer 1 year, 3 months ago

which means A E

upvoted 1 times

Question #106 *Topic 1*

You work for a large organization where each business unit has thousands of users. You need to delegate management of access control permissions to each business unit. You have the following requirements:

➾ Each business unit manages access controls for their own projects.
➾ Each business unit manages access control permissions at scale.
➾ Business units cannot access other business units' projects.
➾ Users lose their access if they move to a different business unit or leave the company.
➾ Users and access control permissions are managed by the on-premises directory service.

What should you do? (Choose two.)

A. Use VPC Service Controls to create perimeters around each business unit's project.

B. Organize projects in folders, and assign permissions to Google groups at the folder level.

C. Group business units based on Organization Units (OUs) and manage permissions based on OUs

D. Create a project naming convention, and use Google's IAM Conditions to manage access based on the prefix of project names.

E. Use Google Cloud Directory Sync to synchronize users and group memberships in Cloud Identity.

Correct Answer: *BE*

pedrojorge Highly Voted 1 year, 7 months ago

Selected Answer: BE

B and E

upvoted 5 times

TheBuckler Highly Voted 1 year, 11 months ago

I will take B & E. Makes sense for the OUs to have their own folders and respective projects under their folders. This will make each OU independent from one another in terms of environments, and will not be able to communicate with one another unless shared VPC/VPC peering utilized.

And E is fairly obvious, as they want to manage their users from on-prem directory, hence GCDS.

upvoted 5 times

tia_gll Most Recent 5 months, 2 weeks ago

Selected Answer: BE

Ans are : B & E

upvoted 1 times

pradoUA 11 months, 1 week ago

Selected Answer: BE

B and E are correct

upvoted 2 times

Rightsaidfred 1 year, 9 months ago

Agreed...B & E

upvoted 3 times

Question #107 *Topic 1*

Your organization recently deployed a new application on Google Kubernetes Engine. You need to deploy a solution to protect the application. The solution has the following requirements:

➾ Scans must run at least once per week
➾ Must be able to detect cross-site scripting vulnerabilities
➾ Must be able to authenticate using Google accounts
Which solution should you use?

A. Google Cloud Armor

B. Web Security Scanner

C. Security Health Analytics

D. Container Threat Detection

Correct Answer: *B*
Reference:
https://cloud.google.com/security-command-center/docs/concepts-web-security-scanner-overview

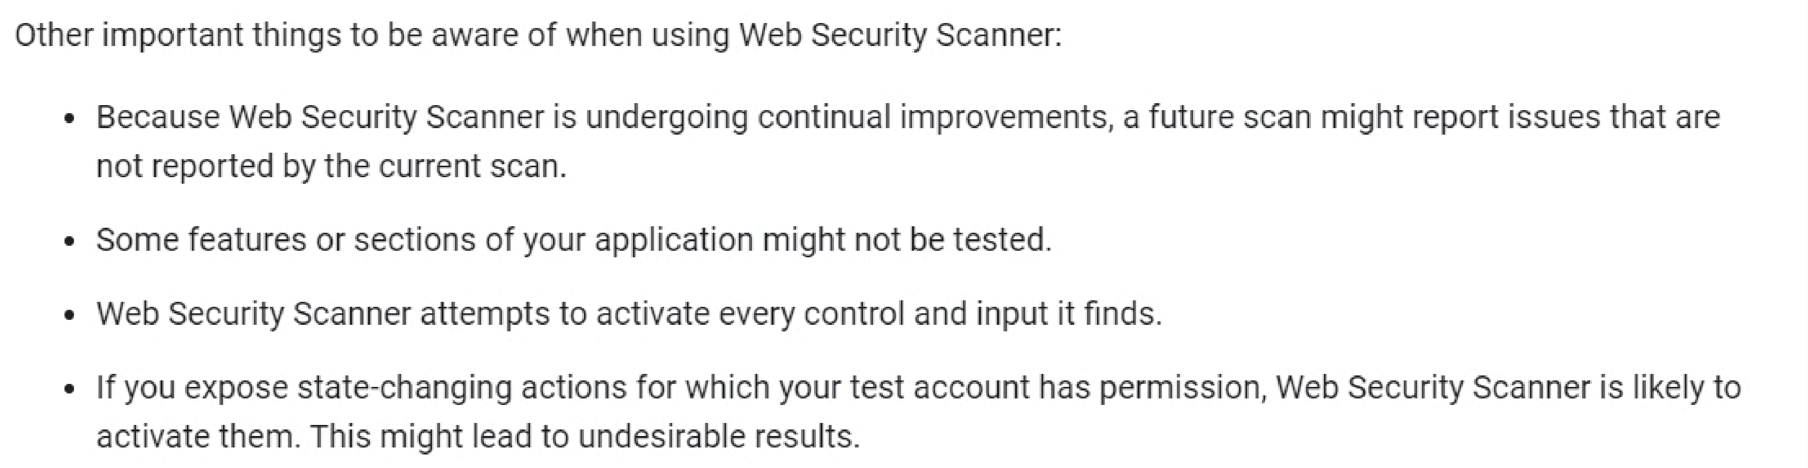

Other important things to be aware of when using Web Security Scanner:

- Because Web Security Scanner is undergoing continual improvements, a future scan might report issues that are not reported by the current scan.
- Some features or sections of your application might not be tested.
- Web Security Scanner attempts to activate every control and input it finds.
- If you expose state-changing actions for which your test account has permission, Web Security Scanner is likely to activate them. This might lead to undesirable results.

Tabayashi Highly Voted 👍 2 years, 4 months ago

Answer is (B).

Web Security Scanner identifies security vulnerabilities in your App Engine, Google Kubernetes Engine (GKE), and Compute Engine web applications.
https://cloud.google.com/security-command-center/docs/concepts-web-security-scanner-overview

upvoted 14 times

AzureDP900 1 year, 10 months ago

Yes, B is right

upvoted 1 times

Alain_Barout2023 Most Recent 🕘 10 months ago

Answer is B.
Web Security Scanner identifies vulnerabilities in web application running in App Engine, Google Kubernetes Engine (GKE), and Compute Engine.
CloudArmor is a WAF solution.

upvoted 3 times

desertlotus1211 8 months ago

Google Cloud Armor can prevent XSS attacks. It has preconfigured rules that can mitigate XSS, broken authentication, and SQL injection. Clou
Armor also has a custom rules
language that includes multiple custom operations.

Could be 'A' as well...

upvoted 1 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: B

B. Web Security Scanner

upvoted 2 times

zellck 1 year, 11 months ago

Selected Answer: B

B is the answer.

upvoted 4 times

Question #108 *Topic 1*

An organization is moving applications to Google Cloud while maintaining a few mission-critical applications on-premises. The organization must transfer the data at a bandwidth of at least 50 Gbps. What should they use to ensure secure continued connectivity between sites?

A. Dedicated Interconnect

B. Cloud Router

C. Cloud VPN

D. Partner Interconnect

Correct Answer: *A*

Reference:

https://cloud.google.com/architecture/migration-to-google-cloud-transferring-your-large-datasets

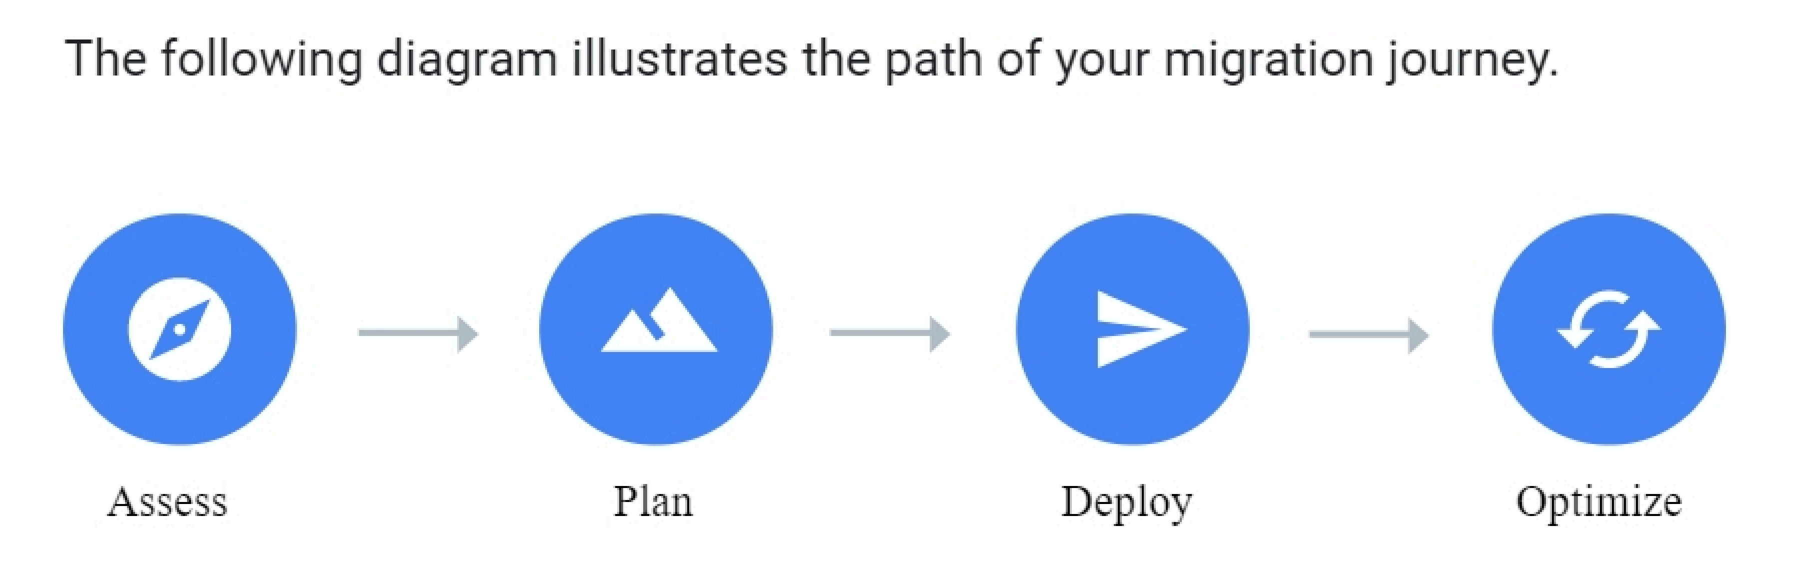

The following diagram illustrates the path of your migration journey.

mouchu Highly Voted 2 years, 3 months ago

Answer = A

upvoted 8 times

[Removed] Highly Voted 1 year, 1 month ago

Selected Answer: A

"A"
I think the keyword here is "at least" 50 Gbps.
Partner interconnect seems to max go up to 50 Gbps but Dedicated Interconnect can guarantee that throughput

https://cloud.google.com/network-connectivity/docs/interconnect/concepts/overview

upvoted 5 times

AzureDP900 Most Recent 1 year, 10 months ago

A is right

upvoted 1 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: A

A. Dedicated Interconnect

upvoted 1 times

zellck 1 year, 11 months ago

Selected Answer: A

A is the answer.

upvoted 1 times

Arturo_Cloud 1 year, 12 months ago

I understand that not all Partner Interconnect connections support 50 Gbps, so I'm going with A) for guaranteed connectivity.

https://cloud.google.com/network-connectivity/docs/interconnect/concepts/overview

upvoted 3 times

Question #109 *Topic 1*

Your organization has had a few recent DDoS attacks. You need to authenticate responses to domain name lookups. Which Google Cloud service should you use?

A. Cloud DNS with DNSSEC

B. Cloud NAT

C. HTTP(S) Load Balancing

D. Google Cloud Armor

Correct Answer: *A*
Reference:
https://developers.google.com/speed/public-dns/faq

Tabayashi Highly Voted 2 years, 4 months ago

Answer is (A).

The Domain Name System Security Extensions (DNSSEC) is a feature of the Domain Name System (DNS) that authenticates responses to domain name lookups. It does not provide privacy protections for those lookups, but prevents attackers from manipulating or poisoning the responses to DNS requests.
https://cloud.google.com/dns/docs/dnssec

upvoted 19 times

AzureDP900 1 year, 10 months ago

Agreed, A is right

upvoted 2 times

Xoxoo Most Recent 11 months, 3 weeks ago

Selected Answer: A

To authenticate responses to domain name lookups and protect your organization from DDoS attacks, you can use Cloud DNS with DNSSEC. DNS Security Extensions (DNSSEC) is a feature of the Domain Name System (DNS) that authenticates responses to domain name lookups and prevents attackers from manipulating or poisoning the responses to DNS requests. Cloud DNS supports DNSSEC and automatically manages the creation and rotation of DNSSEC keys (DNSKEY records) and the signing of zone data with resource record digital signature (RRSIG) records. By enabling DNSSEC in Cloud DNS, you can protect your domains from spoofing and poisoning attacks.

Keyword here is domain name lookup so it must be A.

upvoted 4 times

risc 1 year, 10 months ago

Selected Answer: A

A, as explained by Tabayashi

upvoted 2 times

AwesomeGCP 1 year, 11 months ago

Selected Answer: A

A. Cloud DNS with DNSSEC

upvoted 2 times

zellck 1 year, 11 months ago

Selected Answer: A

A is the answer.

upvoted 2 times

Question #110 *Topic 1*

Your Security team believes that a former employee of your company gained unauthorized access to Google Cloud resources some time in the past 2 months by using a service account key. You need to confirm the unauthorized access and determine the user activity. What should you do?

A. Use Security Health Analytics to determine user activity.

B. Use the Cloud Monitoring console to filter audit logs by user.

C. Use the Cloud Data Loss Prevention API to query logs in Cloud Storage.

D. Use the Logs Explorer to search for user activity.

Correct Answer: *B*

Medofree Highly Voted 2 years, 3 months ago

Selected Answer: D

D.

We use audit logs by searching the Service Account and checking activities in the past 2 months. (the user identity will not be seen since he used the SA identity but we can make correlations based on ip address, working hour, etc.)

upvoted 14 times

AzureDP900 1 year, 10 months ago

D is right, I agree

upvoted 3 times

mikesp Highly Voted 2 years, 3 months ago

Selected Answer: D

B is intended to mislead the public. Cloud Monitoring provides only metrics. To check user activity is necessary to go to Cloud Logging and searcl on Audit Logs.

upvoted 8 times

[Removed] Most Recent 1 year, 1 month ago

Selected Answer: D

"D"

A- Health Analytics - Managed Vulnerability Assessment. Not related.
B- DLP - Filtering/Masing Sensitive Data. Not Related
C- Cloud Monitoring - Perf metrics (e.g. availability). Not related
D- Log Explorer - Log analysis. Related. Great for investigations.

References:
https://cloud.google.com/monitoring
https://cloud.google.com/docs/security/compromised-credentials#look_for_unauthorized_access_and_resources

upvoted 7 times

chickenstealers 1 year, 7 months ago

B is correct answer
https://cloud.google.com/docs/security/compromised-credentials
Monitor for anomalies in service account key usage using Cloud Monitoring.

upvoted 2 times

Sammydp202020 1 year, 6 months ago

Cloud monitoring/logging is a service enabler to capture the logs. Question asks -- How does one check for user activity:

So, the response warranted is D - logs explorer.

https://cloud.google.com/docs/security/compromised-credentials#look_for_unauthorized_access_and_resources

upvoted 1 times

gcpengineer 1 year, 3 months ago

2 months..is long time ti check data access logs

upvoted 1 times

zellck 1 year, 11 months ago

Selected Answer: D

D is the answer.

upvoted 1 times

mT3 2 years, 3 months ago

Selected Answer: B

Correct. Answer is (B).
Investigate the potentially unauthorized activity and restore the account.
Ref.https://support.google.com/a/answer/2984349

upvoted 3 times

Question #111 *Topic 1*

Your company requires the security and network engineering teams to identify all network anomalies within and across VPCs, internal traffic from VMs to VMs, traffic between end locations on the internet and VMs, and traffic between VMs to Google Cloud services in production. Which method should you use?

A. Define an organization policy constraint.

B. Configure packet mirroring policies.

C. Enable VPC Flow Logs on the subnet.

D. Monitor and analyze Cloud Audit Logs.

Correct Answer: *C*

Reference:

https://cloud.google.com/architecture/best-practices-vpc-design

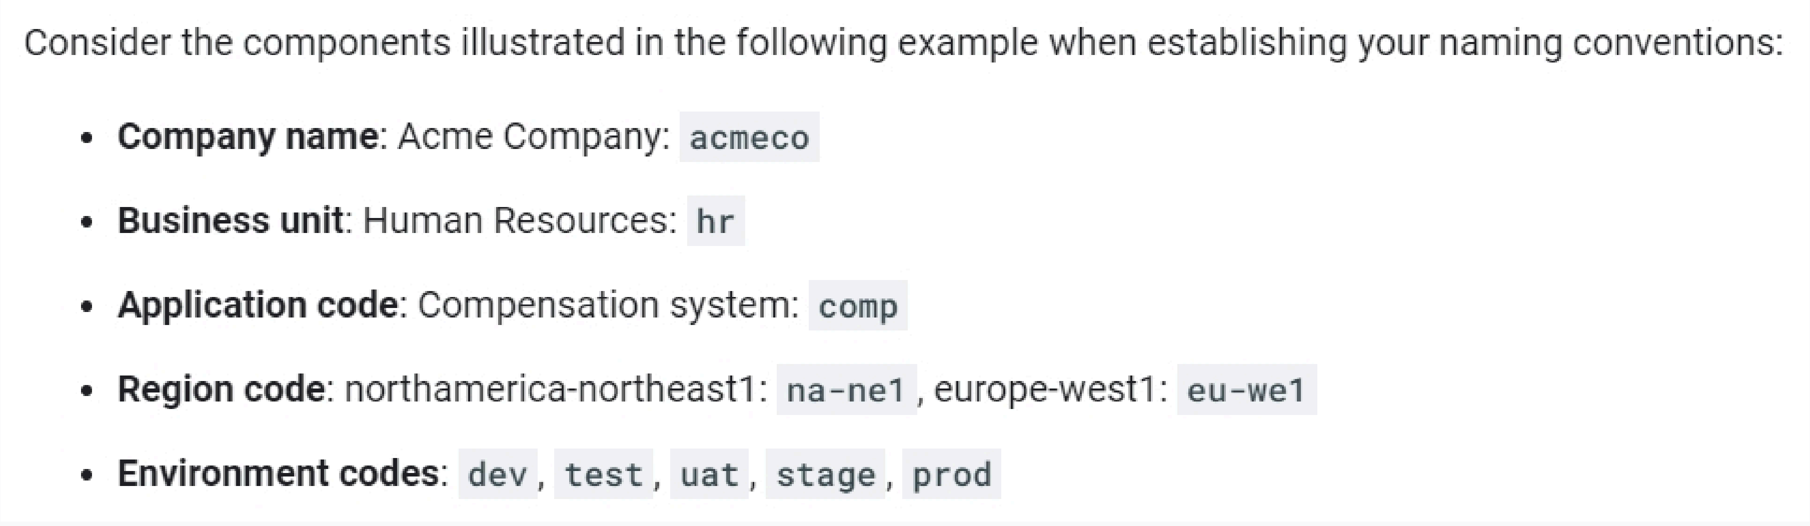

Consider the components illustrated in the following example when establishing your naming conventions:

- **Company name**: Acme Company: `acmeco`
- **Business unit**: Human Resources: `hr`
- **Application code**: Compensation system: `comp`
- **Region code**: northamerica-northeast1: `na-ne1`, europe-west1: `eu-we1`
- **Environment codes**: `dev`, `test`, `uat`, `stage`, `prod`

Tabayashi Highly Voted 2 years, 4 months ago

I think the answer is (C).

VPC Flow Logs samples each VM's TCP, UDP, ICMP, ESP, and GRE flows. Both inbound and outbound flows are sampled. These flows can be between the VM and another VM, a host in your on-premises data center, a Google service, or a host on the internet.
https://cloud.google.com/vpc/docs/flow-logs

upvoted 13 times

hybridpro Highly Voted 2 years, 2 months ago

B should be the answer. For detecting network anomalies, you need to have payload and header data as well to be effective. Besides C is saying t enable VPC flow logs on a subnet which won't serve our purpose either.

upvoted 8 times

dija123 Most Recent 6 months ago

Selected Answer: B

Backet mirroring policies allow you to mirror all traffic passing through a specific network interface or VPC route to a designated destination (e.g., another VM, a Cloud Storage bucket). This captured traffic can then be analyzed by security and network engineers using tools like Suricata or Security Command Center for advanced anomaly detection. This approach provides the necessary level of detail and flexibility for identifying anomalies across all the mentioned traffic types

upvoted 1 times

b6f53d8 8 months ago

C is only for subnet, and we need control in many VPCs, so I prefer B

upvoted 1 times

[Removed] 8 months, 2 weeks ago

Selected Answer: C

C - we need more than just the VMs here.

upvoted 1 times

sebG35 9 months ago

The answer is C. The needs is identify all network anomalies within and across VPCs, internal traffic from VMs to VMs ...

B- Does not meet all needs. It is limited to the VM and don't cover the needs : across VPCs
https://cloud.google.com/vpc/docs/packet-mirroring?hl=en

C- Cover all needs
https://cloud.google.com/vpc/docs/flow-logs?hl=en

upvoted 1 times