



Vendor: Google

Exam Code: Professional-Cloud-Architect

Exam Name: Google Certified Professional - Cloud Architect
(GCP)

Version: DEMO

Case Study: 1 - Mountkirk Games A
Case Study: 2 - TerramEarth A
Case Study: 3 - JencoMart
Case Study: 4 - Dress4Win A
Case Study: 5 - Dress4win B
Case Study: 6 - TerramEarth B
Case Study: 7 - Mountkirk Games B
Case Study: 8 - Mountkirk Games C
Case Study: 9 - Helicopter Racing League
Case Study: 10 - EHR Healthcare
Case Study: 11 - TerramEarth C

QUESTION 1

Your organization requires that metrics from all applications be retained for 5 years for future analysis in possible legal proceedings. Which approach should you use?

- A. Grant the security team access to the logs in each Project.
- B. Configure Stackdriver Monitoring for all Projects, and export to BigQuery.
- C. Configure Stackdriver Monitoring for all Projects with the default retention policies.
- D. Configure Stackdriver Monitoring for all Projects, and export to Google Cloud Storage.

Answer: D

Explanation:

For archival purposes, Customer should use Cloud Storage. BigQuery is a datawarehouse, and could eventually import data from Cloud Storage if necessary.

References:

<https://cloud.google.com/monitoring/api/v3/metrics>

<https://cloud.google.com/stackdriver/>

QUESTION 2

Your company has decided to build a backup replica of their on-premises user authentication PostgreSQL database on Google Cloud Platform. The database is 4 TB, and large updates are frequent. Replication requires private address space communication. Which networking approach should you use?

- A. Google Cloud Dedicated Interconnect
- B. Google Cloud VPN connected to the data center network
- C. A NAT and TLS translation gateway installed on-premises
- D. A Google Compute Engine instance with a VPN server installed connected to the data center network

Answer: A

Explanation:

Google Cloud Dedicated Interconnect provides direct physical connections and RFC 1918 communication between your on-premises network and Google's network. Dedicated Interconnect enables you to transfer large amounts of data between networks, which can be more cost effective than purchasing additional bandwidth over the public Internet or using VPN tunnels.

Benefits:

- Traffic between your on-premises network and your VPC network doesn't traverse the public Internet. Traffic traverses a dedicated connection with fewer hops, meaning there are less points of failure where traffic might get dropped or disrupted.

- Your VPC network's internal (RFC 1918) IP addresses are directly accessible from your on-premises network. You don't need to use a NAT device or VPN tunnel to reach internal IP addresses. Currently, you can only reach internal IP addresses over a dedicated connection. To reach Google external IP addresses, you must use a separate connection.
- You can scale your connection to Google based on your needs. Connection capacity is delivered over one or more 10 Gbps Ethernet connections, with a maximum of eight connections (80 Gbps total per interconnect).
- The cost of egress traffic from your VPC network to your on-premises network is reduced. A dedicated connection is generally the least expensive method if you have a high-volume of traffic to and from Google's network.

References: <https://cloud.google.com/interconnect/docs/details/dedicated>

QUESTION 3

Your company is forecasting a sharp increase in the number and size of Apache Spark and Hadoop jobs being run on your local datacenter. You want to utilize the cloud to help you scale this upcoming demand with the least amount of operations work and code change. Which product should you use?

- A. Google Cloud Dataflow
- B. Google Cloud Dataproc
- C. Google Compute Engine
- D. Google Container Engine

Answer: B

Explanation:

Google Cloud Dataproc is a fast, easy-to-use, low-cost and fully managed service that lets you run the Apache Spark and Apache Hadoop ecosystem on Google Cloud Platform. Cloud Dataproc provisions big or small clusters rapidly, supports many popular job types, and is integrated with other Google Cloud Platform services, such as Google Cloud Storage and Stackdriver Logging, thus helping you reduce TCO.

References: <https://cloud.google.com/dataproc/docs/resources/faq>

QUESTION 4

Your company's test suite is a custom C++ application that runs tests throughout each day on Linux virtual machines. The full test suite takes several hours to complete, running on a limited number of on-premises servers reserved for testing. Your company wants to move the testing infrastructure to the cloud, to reduce the amount of time it takes to fully test a change to the system, while changing the tests as little as possible. Which cloud infrastructure should you recommend?

- A. Google Compute Engine unmanaged instance groups and Network Load Balancer
- B. Google Compute Engine managed instance groups with auto-scaling
- C. Google Cloud Dataproc to run Apache Hadoop jobs to process each test
- D. Google App Engine with Google Stackdriver for logging

Answer: B

Explanation:

Google Compute Engine enables users to launch virtual machines (VMs) on demand. VMs can be launched from the standard images or custom images created by users.

Managed instance groups offer autoscaling capabilities that allow you to automatically add or remove instances from a managed instance group based on increases or decreases in load. Autoscaling helps your applications gracefully handle increases in traffic and reduces cost when the need for resources is lower.

Incorrect Answers:

B: There is no mention of incoming IP data traffic for the custom C++ applications.

C: Apache Hadoop is not fit for testing C++ applications. Apache Hadoop is an open-source software framework used for distributed storage and processing of datasets of big data using the MapReduce programming model.

D: Google App Engine is intended to be used for web applications.

Google App Engine (often referred to as GAE or simply App Engine) is a web framework and cloud computing platform for developing and hosting web applications in Google-managed data centers.

References: <https://cloud.google.com/compute/docs/autoscaler/>

QUESTION 5

Your company just finished a rapid lift and shift to Google Compute Engine for your compute needs. You have another 9 months to design and deploy a more cloud-native solution. Specifically, you want a system that is no-ops and auto-scaling. Which two compute products should you choose? Choose 2 answers

- A. Compute Engine with containers
- B. Google Container Engine with containers
- C. Google App Engine Standard Environment
- D. Compute Engine with custom instance types
- E. Compute Engine with managed instance groups

Answer: BC

Explanation:

B: With Container Engine, Google will automatically deploy your cluster for you, update, patch, secure the nodes.

Kubernetes Engine's cluster autoscaler automatically resizes clusters based on the demands of the workloads you want to run.

C: Solutions like Datastore, BigQuery, AppEngine, etc are truly NoOps.

App Engine by default scales the number of instances running up and down to match the load, thus providing consistent performance for your app at all times while minimizing idle instances and thus reducing cost.

Note: At a high level, NoOps means that there is no infrastructure to build out and manage during usage of the platform. Typically, the compromise you make with NoOps is that you lose control of the underlying infrastructure.

References: <https://www.quora.com/How-well-does-Google-Container-Engine-support-Google-Cloud-Platform%E2%80%99s-NoOps-claim>

QUESTION 6

A development manager is building a new application He asks you to review his requirements and identify what cloud technologies he can use to meet them.

The application must:

1. Be based on open-source technology for cloud portability
2. Dynamically scale compute capacity based on demand
3. Support continuous software delivery
4. Run multiple segregated copies of the same application stack
5. Deploy application bundles using dynamic templates
6. Route network traffic to specific services based on URL

Which combination of technologies will meet all of his requirements?

- A. Google Container Engine, Jenkins, and Helm
- B. Google Container Engine and Cloud Load Balancing
- C. Google Compute Engine and Cloud Deployment Manager
- D. Google Compute Engine, Jenkins, and Cloud Load Balancing

Answer: D

Explanation:

Jenkins is an open-source automation server that lets you flexibly orchestrate your build, test, and deployment pipelines. Kubernetes Engine is a hosted version of Kubernetes, a powerful cluster manager and orchestration system for containers.

When you need to set up a continuous delivery (CD) pipeline, deploying Jenkins on Kubernetes Engine provides important benefits over a standard VM-based deployment

Incorrect Answers:

A: Helm is a tool for managing Kubernetes charts. Charts are packages of pre-configured Kubernetes resources.

Use Helm to:

- Find and use popular software packaged as Kubernetes charts
- Share your own applications as Kubernetes charts
- Create reproducible builds of your Kubernetes applications
- Intelligently manage your Kubernetes manifest files
- Manage releases of Helm packages

References: <https://cloud.google.com/solutions/jenkins-on-kubernetes-engine>

QUESTION 7

You have an application deployed on Kubernetes Engine using a Deployment named echo-deployment. The deployment is exposed using a Service called echo-service. You need to perform an update to the application with minimal downtime to the application.

What should you do?

- A. Use `kubectl set image deployment/echo-deployment <new-image>`
- B. Use the rolling update functionality of the Instance Group behind the Kubernetes cluster
- C. Update the deployment yaml file with the new container image.
Use `kubectl delete deployment/echo-deployment` and `kubectl create -f <yaml-file>`
- D. Update the service yaml file with the new container image.
Use `kubectl delete service/echo-service` and `kubectl create -f <yaml-file>`

Answer: A

Explanation:

Deployment ensures that only a certain number of Pods are down while they are being updated. By default, it ensures that at least 75% of the desired number of Pods are up (25% max unavailable).

Deployment also ensures that only a certain number of Pods are created above the desired number of Pods. By default, it ensures that at most 125% of the desired number of Pods are up (25% max surge).

<https://kubernetes.io/docs/concepts/workloads/controllers/deployment/#updating-a-deployment>

QUESTION 8

Your company is using BigQuery as its enterprise data warehouse. Data is distributed over several Google Cloud projects. All queries on BigQuery need to be billed on a single project. You want to make sure that no query costs are incurred on the projects that contain the data. Users should be able to query the datasets, but not edit them.

How should you configure users' access roles?

- A. Add all users to a group. Grant the group the role of BigQuery user on the billing project and BigQuery dataViewer on the projects that contain the data.
- B. Add all users to a group. Grant the group the roles of BigQuery dataViewer on the billing project and BigQuery user on the projects that contain the data.
- C. Add all users to a group. Grant the group the roles of BigQuery jobUser on the billing project and BigQuery dataViewer on the projects that contain the data.
- D. Add all users to a group. Grant the group the roles of BigQuery dataViewer on the billing project and BigQuery jobUser on the projects that contain the data.

Answer: C

Explanation:

BigQuery User: (roles/bigquery.user)

When applied to a dataset, this role provides the ability to read the dataset's metadata and list tables in the dataset.

When applied to a project, this role also provides the ability to run jobs, including queries, within the project. A principal with this role can enumerate their own jobs, cancel their own jobs, and enumerate datasets within a project. Additionally, allows the creation of new datasets within the project; the creator is granted the BigQuery Data Owner role(roles/bigquery.dataOwner) on these new datasets.

Lowest-level resources where you can grant this role: Dataset

BigQuery Job User: (roles/bigquery.jobUser)

Provides permissions to run jobs, including queries, within the project.

Lowest-level resources where you can grant this role: Project

Source: <https://cloud.google.com/bigquery/docs/access-control>

QUESTION 9

You have developed an application using Cloud ML Engine that recognizes famous paintings from uploaded images. You want to test the application and allow specific people to upload images for the next 24 hours. Not all users have a Google Account. How should you have users upload images?

- A. Have users upload the images to Cloud Storage. Protect the bucket with a password that expires after 24 hours.
- B. Have users upload the images to Cloud Storage using a signed URL that expires after 24 hours.
- C. Create an App Engine web application where users can upload images. Configure App Engine to

- disable the application after 24 hours. Authenticate users via Cloud Identity.
- D. Create an App Engine web application where users can upload images for the next 24 hours. Authenticate users via Cloud Identity.

Answer: B

Explanation:

When should you use a signed URL? In some scenarios, you might not want to require your users to have a Google account in order to access Cloud Storage. "Signed URLs contain authentication information in their query string, allowing users without credentials to perform specific actions on a resource".

<https://cloud.google.com/storage/docs/access-control/signed-urls>

QUESTION 10

Your company operates nationally and plans to use GCP for multiple batch workloads, including some that are not time-critical. You also need to use GCP services that are HIPAA-certified and manage service costs.

How should you design to meet Google best practices?

- A. Provisioning preemptible VMs to reduce cost. Discontinue use of all GCP services and APIs that are not HIPAA-compliant.
- B. Provisioning preemptible VMs to reduce cost. Disable and then discontinue use of all GCP and APIs that are not HIPAA-compliant.
- C. Provision standard VMs in the same region to reduce cost. Discontinue use of all GCP services and APIs that are not HIPAA-compliant.
- D. Provision standard VMs to the same region to reduce cost. Disable and then discontinue use of all GCP services and APIs that are not HIPAA-compliant.

Answer: B

Explanation:

<https://cloud.google.com/compute/docs/instances/preemptible>

If your apps are fault-tolerant and can withstand possible instance preemptions, then preemptible instances can reduce your Compute Engine costs significantly. For example, batch processing jobs can run on preemptible instances. If some of those instances stop during processing, the job slows but does not completely stop. Preemptible instances complete your batch processing tasks without placing additional workload on your existing instances and without requiring you to pay full price for additional normal instances.

QUESTION 11

You have been engaged by your client to lead the migration of their application infrastructure to GCP. One of their current problems is that the on-premises high performance SAN is requiring frequent and expensive upgrades to keep up with the variety of workloads that are identified as follows: 20TB of log archives retained for legal reasons; 500 GB of VM boot/data volumes and templates; 500 GB of image thumbnails; 200 GB of customer session state data that allows customers to restart sessions even if off-line for several days.

Which of the following best reflects your recommendations for a cost-effective storage allocation?

- A. Local SSD for customer session state data. Lifecycle-managed Cloud Storage for log archives, thumbnails, and VM boot/data volumes.
- B. Memcache backed by Cloud Datastore for the customer session state data. Lifecycle-managed Cloud Storage for log archives, thumbnails, and VM boot/data volumes.

- C. Memcache backed by Cloud SQL for customer session state data. Assorted local SSD-backed instances for VM boot/data volumes. Cloud Storage for log archives and thumbnails.
- D. Memcache backed by Persistent Disk SSD storage for customer session state data. Assorted local SSD-backed instances for VM boot/data volumes. Cloud Storage for log archives and thumbnails.

Answer: B

Explanation:

Memcache backed by Cloud Datastore

<https://cloud.google.com/appengine/docs/standard/python/memcache>

Compute Engine image can be stored in Cloud Storage

<https://cloud.google.com/solutions/image-management-best-practices>

After the complete sequence of bytes from the disk are written to the file, the file is archived using the tar format and then compressed using the GZIP format. You can then upload the resulting *.tar.gz file to Cloud Storage and register it as an image in Compute Engine.

QUESTION 12

You are using Cloud CDN to deliver static HTTP(S) website content hosted on a Compute Engine instance group. You want to improve the cache hit ratio.

What should you do?

- A. Customize the cache keys to omit the protocol from the key.
- B. Shorten the expiration time of the cached objects.
- C. Make sure the HTTP(S) header "Cache-Region" points to the closest region of your users.
- D. Replicate the static content in a Cloud Storage bucket. Point CloudCDN toward a load balancer on that bucket.

Answer: A

Explanation:

A logo needs to be cached whether displayed through HTTP or HTTPS. When you customize the cache keys for the backend service that holds the logo, clear the Protocol checkbox so that requests through HTTP and HTTPS count as matches for the logo's cache entry.

https://cloud.google.com/cdn/docs/best-practices#using_custom_cache_keys_to_improve_cache_hit_ratio

QUESTION 13

All compute Engine instances in your VPC should be able to connect to an Active Directory server on specific ports. Any other traffic emerging from your instances is not allowed. You want to enforce this using VPC firewall rules.

How should you configure the firewall rules?

- A. Create an egress rule with priority 1000 to deny all traffic for all instances. Create another egress rule with priority 100 to allow the Active Directory traffic for all instances.
- B. Create an egress rule with priority 100 to deny all traffic for all instances. Create another egress rule with priority 1000 to allow the Active Directory traffic for all instances.
- C. Create an egress rule with priority 1000 to allow the Active Directory traffic. Rely on the implied deny egress rule with priority 100 to block all traffic for all instances.
- D. Create an egress rule with priority 100 to allow the Active Directory traffic. Rely on the implied deny egress rule with priority 1000 to block all traffic for all instances.

Answer: A

Explanation:

https://cloud.google.com/vpc/docs/firewalls#default_firewall_rules

Every VPC network has two implied firewall rules. These rules exist, but are not shown in the Cloud Console:

The implied 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, except for traffic blocked by GCP. Outbound access may be restricted by a higher priority firewall rule. Internet access is allowed if no other firewall rules deny outbound traffic and if the instance has an external IP address or uses a NAT instance. Refer to Internet access requirements for more details.

The implied 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 traffic to them. Incoming access may be allowed by a higher priority rule. Note that the default network includes some additional rules that override this one, allowing certain types of incoming traffic.

QUESTION 14

You are developing a globally scaled frontend for a legacy streaming backend data API. This API expects events in strict chronological order with no repeat data for proper processing.

Which products should you deploy to ensure guaranteed-once FIFO (first-in, first-out) delivery of data?

- A. Cloud Pub/Sub alone
- B. Cloud Pub/Sub to Cloud DataFlow
- C. Cloud Pub/Sub to Stackdriver
- D. Cloud Pub/Sub to Cloud SQL

Answer: B

Explanation:

CloudSQL is only for storage, to get the messages in order you need timestamp processed in dataflow to arrange them before putting it in any storage volume. The system described is not querying a db it is expecting a stream of messages only dataflow can correct the order. ACID has no value here because the db is not being queried. You'll not find any documentation on pub/sub order being corrected with a db. See notes below on pub/sub and dataflow using timestamps and windows to ensure order.

<https://cloud.google.com/pubsub/docs/pubsub-dataflow>

QUESTION 15

A recent audit revealed that a new network was created in your GCP project. In this network, a GCE instance has an SSH port open to the world. You want to discover this network's origin.

What should you do?

- A. Search for Create VM entry in the Stackdriver alerting console
- B. Navigate to the Activity page in the Home section. Set category to Data Access and search for Create VM entry
- C. In the Logging section of the console, specify GCE Network as the logging section. Search for the Create Insert entry
- D. Connect to the GCE instance using project SSH keys. Identify previous logins in system logs, and match these with the project owners list

Answer: C

Explanation:

Incorrect Answers:

A: To use the Stackdriver alerting console we must first set up alerting policies.

B: Data access logs only contain read-only operations.

Audit logs help you determine who did what, where, and when.

Cloud Audit Logging returns two types of logs:

Admin activity logs

Data access logs: Contains log entries for operations that perform read-only operations do not modify any data, such as get, list, and aggregated list methods.

QUESTION 16

A production database virtual machine on Google Compute Engine has an ext4-formatted persistent disk for data files. The database is about to run out of storage space.

How can you remediate the problem with the least amount of downtime?

- A. In the Cloud Platform Console, increase the size of the persistent disk and use the `resize2fs` command in Linux.
- B. Shut down the virtual machine, use the Cloud Platform Console to increase the persistent disk size, then restart the virtual machine
- C. In the Cloud Platform Console, increase the size of the persistent disk and verify the new space is ready to use with the `fdisk` command in Linux
- D. In the Cloud Platform Console, create a new persistent disk attached to the virtual machine, format and mount it, and configure the database service to move the files to the new disk
- E. In the Cloud Platform Console, create a snapshot of the persistent disk restore the snapshot to a new larger disk, unmount the old disk, mount the new disk and restart the database service

Answer: A

Explanation:

On Linux instances, connect to your instance and manually resize your partitions and file systems to use the additional disk space that you added.

Extend the file system on the disk or the partition to use the added space. If you grew a partition on your disk, specify the partition. If your disk does not have a partition table, specify only the disk ID.

```
sudo resize2fs /dev/[DISK_ID][PARTITION_NUMBER]
```

where [DISK_ID] is the device name and [PARTITION_NUMBER] is the partition number for the device where you are resizing the file system.

References: <https://cloud.google.com/compute/docs/disks/add-persistent-disk>

QUESTION 17

Your company has a Google Cloud project that uses BigQuery for data warehousing. They have a VPN tunnel between the on-premises environment and Google Cloud that is configured with Cloud VPN. The security team wants to avoid data exfiltration by malicious insiders, compromised code, and accidental oversharing. What should they do?

- A. Configure Private Google Access for on-premises only.
- B. Perform the following tasks:
 - 1. Create a service account.
 - 2. Give the BigQuery JobUser role and Storage Reader role to the service account.
 - 3. Remove all other IAM access from the project.
- C. Configure VPC Service Controls and configure Private Google Access.

D. Configure Private Google Access.

Answer: C

Explanation:

Data exfiltration by malicious insiders or compromised code: VPC Service Controls complements network egress controls by preventing clients within those networks from accessing the resources of Google-managed services outside the perimeter.

<https://cloud.google.com/vpc-service-controls/docs/overview>

QUESTION 18

You want to allow your operations team to store logs from all the production projects in your Organization, without including logs from other projects. All of the production projects are contained in a folder. You want to ensure that all logs for existing and new production projects are captured automatically. What should you do?

- A. Create an aggregated export on the Production folder. Set the log sink to be a Cloud Storage bucket in an operations project.
- B. Create an aggregated export on the Organization resource. Set the log sink to be a Cloud Storage bucket in an operations project.
- C. Create log exports in the production projects. Set the log sinks to be a Cloud Storage bucket in an operations project.
- D. Create log exports in the production projects. Set the log sinks to be BigQuery datasets in the production projects, and grant IAM access to the operations team to run queries on the datasets.

Answer: A

Explanation:

https://cloud.google.com/logging/docs/export/aggregated_sinks

Aggregated sinks combine and route log entries from the Google Cloud resources contained by an organization or folder.

QUESTION 19

Your company has a Google Cloud project that uses BigQuery for data warehousing. There are some tables that contain personally identifiable information (PII). Only the compliance team may access the PII. The other information in the tables must be available to the data science team. You want to minimize cost and the time it takes to assign appropriate access to the tables. What should you do?

- A.
 1. From the dataset where you have the source data, create views of tables that you want to share, excluding PII.
 2. Assign an appropriate project-level IAM role to the members of the data science team.
 3. Assign access controls to the dataset that contains the view.
- B.
 1. From the dataset where you have the source data, create materialized views of tables that you want to share, excluding PII.
 2. Assign an appropriate project-level IAM role to the members of the data science team.
 3. Assign access controls to the dataset that contains the view.
- C.
 1. Create a dataset for the data science team.
 2. Create views of tables that you want to share, excluding PII.
 3. Assign an appropriate project-level IAM role to the members of the data science team.
 4. Assign access controls to the dataset that contains the view.
 5. Authorize the view to access the source dataset.
- D.
 1. Create a dataset for the data science team.
 2. Create materialized views of tables that you want to share, excluding PII.

3. Assign an appropriate project-level IAM role to the members of the data science team.
4. Assign access controls to the dataset that contains the view.
5. Authorize the view to access the source dataset.

Answer: C

Explanation:

You need to create a new dataset based on a query from the source dataset first before you can create the authorised view.

https://cloud.google.com/bigquery/docs/share-access-views#create_a_dataset_where_you_can_store_your_view

After creating your source dataset, you create a new, separate dataset to store the authorized view that you share with your data analysts. In a later step, you grant the authorized view access to the data in the source dataset. Your data analysts then have access to the authorized view, but not direct access to the source data.

QUESTION 20

Your company and one of its partners each have a Google Cloud project in separate organizations.

Your company's project (prj-a) runs in Virtual Private Cloud (vpc-a). The partner's project (prj-b) runs in vpc-b. There are two instances running on vpc-a and one instance running on vpc-b. Subnets defined in both VPCs are not overlapping. You need to ensure that all instances communicate with each other via internal IPs, minimizing latency and maximizing throughput. What should you do?

- A. Set up a network peering between vpc-a and vpc-b
- B. Set up a VPN between vpc-a and vpc-b using Cloud VPN
- C. Configure IAP TCP forwarding on the instance in vpc-b, and then launch the following gcloud command from one of the instances in vpc-a `gcloud compute start-iap-tunnel INSTANCE_NAME_IN_VPC_8 22 \ --local-host-port=localhost:22`
- D.
 1. Create an additional instance in vpc-a.
 2. Create an additional instance in vpc-b.
 3. Install OpenVPN in newly created instances.
 4. Configure a VPN tunnel between vpc-a and vpc-b with the help of OpenVPN.

Answer: A

Explanation:

<https://cloud.google.com/vpc/docs/vpc-peering>

Google Cloud VPC Network Peering allows internal IP address connectivity across two Virtual Private Cloud (VPC) networks regardless of whether they belong to the same project or the same organization.

Thank You for Trying Our Product

Lead2pass Certification Exam Features:

- ★ More than **99,900** Satisfied Customers Worldwide.
- ★ Average **99.9%** Success Rate.
- ★ **Free Update** to match latest and real exam scenarios.
- ★ **Instant Download** Access! No Setup required.
- ★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.
- ★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad**.
- ★ **100%** Guaranteed Success or **100%** Money Back Guarantee.
- ★ **Fast**, helpful support **24x7**.



View list of all certification exams: <http://www.lead2pass.com/all-products.html>



Microsoft



ORACLE



JUNIPER
NETWORKS



EMC²
where information lives

10% Discount Coupon Code: ASTR14