



**Vendor:** Amazon

**Exam Code:** SAA-C03

**Exam Name:** AWS Certified Solutions Architect - Associate  
(SAA-C03) Exam

**Version:** DEMO

### QUESTION 1

A solutions architect is designing a new hybrid architecture to extend a company's on-premises infrastructure to AWS. The company requires a highly available connection with consistent low latency to an AWS Region. The company needs to minimize costs and is willing to accept slower traffic if the primary connection fails.

What should the solutions architect do to meet these requirements?

- A. Provision an AWS Direct Connect connection to a Region.  
Provision a VPN connection as a backup if the primary Direct Connect connection fails.
- B. Provision a VPN tunnel connection to a Region for private connectivity.  
Provision a second VPN tunnel for private connectivity and as a backup if the primary VPN connection fails.
- C. Provision an AWS Direct Connect connection to a Region.  
Provision a second Direct Connect connection to the same Region as a backup if the primary Direct Connect connection fails.
- D. Provision an AWS Direct Connect connection to a Region.  
Use the Direct Connect failover attribute from the AWS CLI to automatically create a backup connection if the primary Direct Connect connection fails.

**Answer:** A

#### **Explanation:**

"In some cases, this connection alone is not enough. It is always better to guarantee a fallback connection as the backup of DX. There are several options, but implementing it with an AWS Site-To-Site VPN is a real cost-effective solution that can be exploited to reduce costs or, in the meantime, wait for the setup of a second DX."

<https://www.proud2becloud.com/hybrid-cloud-networking-backup-aws-direct-connect-network-connection-with-aws-site-to-site-vpn/>

### QUESTION 2

A solutions architect is designing a two-tier web application. The application consists of a public-facing web tier hosted on Amazon EC2 in public subnets. The database tier consists of Microsoft SQL Server running on Amazon EC2 in a private subnet. Security is a high priority for the company.

How should security groups be configured in this situation? (Select TWO )

- A. Configure the security group for the web tier to allow inbound traffic on port 443 from 0.0.0.0/0.
- B. Configure the security group for the web tier to allow outbound traffic on port 443 from 0.0.0.0/0.
- C. Configure the security group for the database tier to allow inbound traffic on port 1433 from the security group for the web tier.
- D. Configure the security group for the database tier to allow outbound traffic on ports 443 and 1433 to the security group for the web tier.
- E. Configure the security group for the database tier to allow inbound traffic on ports 443 and 1433 from the security group for the web tier.

**Answer:** AC

#### **Explanation:**

"Security groups create an outbound rule for every inbound rule." Not completely right. Stateful does NOT mean that if you create an inbound (or outbound) rule, it will create an outbound (or inbound) rule. What it does mean is: suppose you create an inbound rule on port 443 for the X ip. When a request enters on port 443 from X ip, it will allow traffic out for that request on the port 443. However, if you look at the outbound rules, there will not be any outbound rule on port 443 unless explicitly create it. In ACLs, which are stateless, you would have to create an inbound rule to allow incoming requests and an outbound rule to allow your application responds to those

incoming requests.

[https://docs.aws.amazon.com/vpc/latest/userguide/VPC\\_SecurityGroups.html#SecurityGroupRules](https://docs.aws.amazon.com/vpc/latest/userguide/VPC_SecurityGroups.html#SecurityGroupRules)

### QUESTION 3

A company receives 10 TB of instrumentation data each day from several machines located at a single factory. The data consists of JSON files stored on a storage area network (SAN) in an on-premises data center located within the factory. The company wants to send this data to Amazon S3 where it can be accessed by several additional systems that provide critical near-real-time analytics.

A secure transfer is important because the data is considered sensitive.

Which solution offers the MOST reliable data transfer?

- A. AWS DataSync over public internet
- B. AWS DataSync over AWS Direct Connect
- C. AWS Database Migration Service (AWS DMS) over public internet
- D. AWS Database Migration Service (AWS DMS) over AWS Direct Connect

**Answer: B**

#### **Explanation:**

These are some of the main use cases for AWS DataSync:

- Data migration
- Move active datasets rapidly over the network into Amazon S3, Amazon EFS, or FSx for Windows File Server.

DataSync includes automatic encryption and data integrity validation to help make sure that your data arrives securely, intact, and ready to use.

DataSync includes encryption and integrity validation to help make sure your data arrives securely, intact, and ready to use.

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

### QUESTION 4

A solutions architect is designing the cloud architecture for a new application being deployed on AWS. The process should run in parallel while adding and removing application nodes as needed based on the number of jobs to be processed. The processor application is stateless. The solutions architect must ensure that the application is loosely coupled and the job items are durably stored.

Which design should the solutions architect use?

- A. Create an Amazon SNS topic to send the jobs that need to be processed.  
Create an Amazon Machine Image (AMI) that consists of the processor application.  
Create a launch configuration that uses the AMI.  
Create an Auto Scaling group using the launch configuration.  
Set the scaling policy for the Auto Scaling group to add and remove nodes based on CPU usage.
- B. Create an Amazon SQS queue to hold the jobs that need to be processed.  
Create an Amazon Machine image (AMI) that consists of the processor application.  
Create a launch configuration that uses the AMI.  
Create an Auto Scaling group using the launch configuration.  
Set the scaling policy for the Auto Scaling group to add and remove nodes based on network usage.
- C. Create an Amazon SQS queue to hold the jobs that needs to be processed.  
Create an Amazon Machine image (AMI) that consists of the processor application.

Create a launch template that uses the AMI.  
Create an Auto Scaling group using the launch template.  
Set the scaling policy for the Auto Scaling group to add and remove nodes based on the number of items in the SQS queue.

- D. Create an Amazon SNS topic to send the jobs that need to be processed.  
Create an Amazon Machine Image (AMI) that consists of the processor application.  
Create a launch template that uses the AMI.  
Create an Auto Scaling group using the launch template.  
Set the scaling policy for the Auto Scaling group to add and remove nodes based on the number of messages published to the SNS topic

**Answer: C**

**Explanation:**

"Create an Amazon SQS queue to hold the jobs that needs to be processed. Create an Amazon EC2 Auto Scaling group for the compute application. Set the scaling policy for the Auto Scaling group to add and remove nodes based on the number of items in the SQS queue"

In this case we need to find a durable and loosely coupled solution for storing jobs. Amazon SQS is ideal for this use case and can be configured to use dynamic scaling based on the number of jobs waiting in the queue. To configure this scaling you can use the backlog per instance metric with the target value being the acceptable backlog per instance to maintain. You can calculate these numbers as follows: Backlog per instance: To calculate your backlog per instance, start with the ApproximateNumberOfMessages queue attribute to determine the length of the SQS queue

#### QUESTION 5

A company's dynamic website is hosted using on-premises servers in the United States. The company is launching its product in Europe, and it wants to optimize site loading times for new European users. The site's backend must remain in the United States. The product is being launched in a few days, and an immediate solution is needed.

What should the solutions architect recommend?

- A. Launch an Amazon EC2 instance in us-east-1 and migrate the site to it.
- B. Move the website to Amazon S3. Use cross-Region replication between Regions.
- C. Use Amazon CloudFront with a custom origin pointing to the on-premises servers.
- D. Use an Amazon Route 53 geo-proximity routing policy pointing to on-premises servers.

**Answer: C**

**Explanation:**

<https://aws.amazon.com/pt/blogs/aws/amazon-cloudfront-support-for-custom-origins/>

You can now create a CloudFront distribution using a custom origin. Each distribution will can point to an S3 or to a custom origin. This could be another storage service, or it could be something more interesting and more dynamic, such as an EC2 instance or even an Elastic Load Balancer

#### QUESTION 6

A company has several web servers that need to frequently access a common Amazon RDS MySQL Multi-AZ DB instance. The company wants a secure method for the web servers to connect to the database while meeting a security requirement to rotate user credentials frequently.

Which solution meets these requirements?

- A. Store the database user credentials in AWS Secrets Manager.

Grant the necessary IAM permissions to allow the web servers to access AWS Secrets Manager.

- B. Store the database user credentials in AWS Systems Manager OpsCenter.  
Grant the necessary IAM permissions to allow the web servers to access OpsCenter.
- C. Store the database user credentials in a secure Amazon S3 bucket.  
Grant the necessary IAM permissions to allow the web servers to retrieve credentials and access the database.
- D. Store the database user credentials in files encrypted with AWS Key Management Service (AWS KMS) on the web server file system. The web server should be able to decrypt the files and access the database.

**Answer: A**

**Explanation:**

AWS Secrets Manager helps you protect secrets needed to access your applications, services, and IT resources. The service enables you to easily rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle.

<https://docs.aws.amazon.com/secretsmanager/latest/userguide/intro.html>

**QUESTION 7**

A company hosts an application on AWS Lambda functions that are invoked by an Amazon API Gateway API. The Lambda functions save customer data to an Amazon Aurora MySQL database. Whenever the company upgrades the database, the Lambda functions fail to establish database connections until the upgrade is complete. The result is that customer data is not recorded for some of the event.

A solutions architect needs to design a solution that stores customer data that is created during database upgrades.

Which solution will meet these requirements?

- A. Provision an Amazon RDS proxy to sit between the Lambda functions and the database. Configure the Lambda functions to connect to the RDS proxy.
- B. Increase the run time of the Lambda functions to the maximum. Create a retry mechanism in the code that stores the customer data in the database.
- C. Persist the customer data to Lambda local storage. Configure new Lambda functions to scan the local storage to save the customer data to the database.
- D. Store the customer data in an Amazon Simple Queue Service (Amazon SQS) FIFO queue. Create a new Lambda function that polls the queue and stores the customer data in the database.

**Answer: D**

**Explanation:**

<https://www.learnaws.org/2020/12/13/aws-rds-proxy-deep-dive/>

RDS proxy can improve application availability in such a situation by waiting for the new database instance to be functional and maintaining any requests received from the application during this time. The end result is that the application is more resilient to issues with the underlying database. This will enable solution to hold data till the time DB comes back to normal. RDS proxy is to optimally utilize the connection between Lambda and DB. Lambda can open multiple connection concurrently which can be taxing on DB compute resources, hence RDS proxy was introduced to manage and leverage these connections efficiently.

**QUESTION 8**

A company is preparing to store confidential data in Amazon S3. For compliance reasons the data must be encrypted at rest. Encryption key usage must be logged for auditing purposes. Keys

must be rotated every year.

Which solution meets these requirements and the MOST operationally efficient?

- A. Server-side encryption with customer-provided keys (SSE-C)
- B. Server-side encryption with Amazon S3 managed keys (SSE-S3)
- C. Server-side encryption with AWS KMS (SSE-KMS) customer master keys (CMKs) with manual rotation
- D. Server-side encryption with AWS KMS (SSE-KMS) customer master keys (CMKs) with automate rotation

**Answer: D**

**Explanation:**

<https://docs.aws.amazon.com/kms/latest/developerguide/rotate-keys.html>

When you enable automatic key rotation for a customer managed key, AWS KMS generates new cryptographic material for the KMS key every year. AWS KMS also saves the KMS key's older cryptographic material in perpetuity so it can be used to decrypt data that the KMS key encrypted. Key rotation in AWS KMS is a cryptographic best practice that is designed to be transparent and easy to use. AWS KMS supports optional automatic key rotation only for customer managed CMKs. Enable and disable key rotation. Automatic key rotation is disabled by default on customer managed CMKs.

When you enable (or re-enable) key rotation, AWS KMS automatically rotates the CMK 365 days after the enable date and every 365 days thereafter.

#### QUESTION 9

A company has a dynamic web application hosted on two Amazon EC2 instances. The company has its own SSL certificate, which is on each instance to perform SSL termination.

There has been an increase in traffic recently, and the operations team determined that SSL encryption and decryption is causing the compute capacity of the web servers to reach their maximum limit.

What should a solutions architect do to increase the application's performance?

- A. Create a new SSL certificate using AWS Certificate Manager (ACM) install the ACM certificate on each instance.
- B. Create an Amazon S3 bucket Migrate the SSL certificate to the S3 bucket. Configure the EC2 instances to reference the bucket for SSL termination.
- C. Create another EC2 instance as a proxy server Migrate the SSL certificate to the new instance and configure it to direct connections to the existing EC2 instances
- D. Import the SSL certificate into AWS Certificate Manager (ACM). Create an Application Load Balancer with an HTTPS listener that uses the SSL certificate from ACM.

**Answer: D**

**Explanation:**

<https://aws.amazon.com/certificate-manager/>:

"With AWS Certificate Manager, you can quickly request a certificate, deploy it on ACM-integrated AWS resources, such as Elastic Load Balancers, Amazon CloudFront distributions, and APIs on API Gateway, and let AWS Certificate Manager handle certificate renewals. It also enables you to create private certificates for your internal resources and manage the certificate lifecycle centrally."

#### QUESTION 10

A company runs its two-tier ecommerce website on AWS. The web tier consists of a load balancer that sends traffic to Amazon EC2 instances. The database tier uses an Amazon RDS DB instance. The EC2 instances and the RDS DB instance should not be exposed to the public internet. The EC2 instances require internet access to complete payment processing of orders through a third-party web service. The application must be highly available.

Which combination of configuration options will meet these requirements? (Choose two.)

- A. Use an Auto Scaling group to launch the EC2 instances in private subnets.  
Deploy an RDS Multi-AZ DB instance in private subnets.
- B. Configure a VPC with two private subnets and two NAT gateways across two Availability Zones.  
Deploy an Application Load Balancer in the private subnets.
- C. Use an Auto Scaling group to launch the EC2 instances in public subnets across two Availability Zones.  
Deploy an RDS Multi-AZ DB instance in private subnets.
- D. Configure a VPC with one public subnet, one private subnet, and two NAT gateways across two Availability Zones.  
Deploy an Application Load Balancer in the public subnet.
- E. Configure a VPC with two public subnets, two private subnets, and two NAT gateways across two Availability Zones.  
Deploy an Application Load Balancer in the public subnets.

**Answer:** AE

**Explanation:**

Before you begin: Decide which two Availability Zones you will use for your EC2 instances. Configure your virtual private cloud (VPC) with at least one public subnet in each of these Availability Zones. These public subnets are used to configure the load balancer. You can launch your EC2 instances in other subnets of these Availability Zones instead.

#### QUESTION 11

A company recently started using Amazon Aurora as the data store for its global ecommerce application.

When large reports are run developers report that the ecommerce application is performing poorly. After reviewing metrics in Amazon CloudWatch, a solutions architect finds that the ReadIOPS and CPU Utilization metrics are spiking when monthly reports run.

What is the MOST cost-effective solution?

- A. Migrate the monthly reporting to Amazon Redshift.
- B. Migrate the monthly reporting to an Aurora Replica
- C. Migrate the Aurora database to a larger instance class
- D. Increase the Provisioned IOPS on the Aurora instance

**Answer:** B

**Explanation:**

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Replication.html>  
#Aurora.Replication.Replicas Aurora Replicas have two main purposes. You can issue queries to them to scale the read operations for your application. You typically do so by connecting to the reader endpoint of the cluster. That way, Aurora can spread the load for read-only connections across as many Aurora Replicas as you have in the cluster. Aurora Replicas also help to increase availability. If the writer instance in a cluster becomes unavailable, Aurora automatically promotes one of the reader instances to take its place as the new writer.

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Overview.html>

### QUESTION 12

A company has a service that produces event data. The company wants to use AWS to process the event data as it is received.

The data is written in a specific order that must be maintained throughout processing.

The company wants to implement a solution that minimizes operational overhead.

How should a solutions architect accomplish this?

- A. Create an Amazon Simple Queue Service (Amazon SQS) FIFO queue to hold messages. Set up an AWS Lambda function to process messages from the queue.
- B. Create an Amazon Simple Notification Service (Amazon SNS) topic to deliver notifications containing payloads to process. Configure an AWS Lambda function as a subscriber.
- C. Create an Amazon Simple Queue Service (Amazon SQS) standard queue to hold messages. Set up an AWS Lambda function to process messages from the queue independently.
- D. Create an Amazon Simple Notification Service (Amazon SNS) topic to deliver notifications containing payloads to process. Configure an Amazon Simple Queue Service (Amazon SQS) queue as a subscriber.

**Answer:** A

#### **Explanation:**

The details are revealed in below url:

<https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/FIFO-queues.html>

FIFO (First-In-First-Out) queues are designed to enhance messaging between applications when the order of operations and events is critical, or where duplicates can't be tolerated. Examples of situations where you might use FIFO queues include the following: To make sure that user-entered commands are run in the right order. To display the correct product price by sending price modifications in the right order. To prevent a student from enrolling in a course before registering for an account.

### QUESTION 13

A large media company hosts a web application on AWS. The company wants to start caching confidential media files so that users around the world will have reliable access to the files. The content is stored in Amazon S3 buckets. The company must deliver the content quickly, regardless of where the requests originate geographically.

Which solution will meet these requirements?

- A. Use AWS DataSync to connect the S3 buckets to the web application.
- B. Deploy AWS Global Accelerator to connect the S3 buckets to the web application.
- C. Deploy Amazon CloudFront to connect the S3 buckets to CloudFront edge servers.
- D. Use Amazon Simple Queue Service (Amazon SQS) to connect the S3 buckets to the web application.

**Answer:** C

#### **Explanation:**

CloudFront uses a local cache to provide the response, AWS Global accelerator proxies requests and connects to the application all the time for the response.

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/private-content-restricting-access-to-s3.html#private-content-granting-permissions-to-oai>



#### QUESTION 14

A company runs a high performance computing (HPC) workload on AWS. The workload required low-latency network performance and high network throughput with tightly coupled node-to-node communication. The Amazon EC2 instances are properly sized for compute and storage capacity, and are launched using default options.

What should a solutions architect propose to improve the performance of the workload?

- A. Choose a cluster placement group while launching Amazon EC2 instances.
- B. Choose dedicated instance tenancy while launching Amazon EC2 instances.
- C. Choose an Elastic Inference accelerator while launching Amazon EC2 instances.
- D. Choose the required capacity reservation while launching Amazon EC2 instances.

**Answer: A**

#### **Explanation:**

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-ec2-placementgroup.html>

A cluster placement group is a logical grouping of instances within a single Availability Zone that benefit from low network latency, high network throughput.

#### QUESTION 15

A company wants to use AWS Systems Manager to manage a fleet of Amazon EC2 instances. According to the company's security requirements, no EC2 instances can have internet access. A solutions architect needs to design network connectivity from the EC2 instances to Systems Manager while fulfilling this security obligation. Which solution will meet these requirements?

- A. Deploy the EC2 instances into a private subnet with no route to the internet.
- B. Configure an interface VPC endpoint for Systems Manager. Update routes to use the endpoint.
- C. Deploy a NAT gateway into a public subnet. Configure private subnets with a default route to the NAT gateway.
- D. Deploy an internet gateway. Configure a network ACL to deny traffic to all destinations except Systems Manager.

**Answer: B**

#### **Explanation:**

VPC Peering connections

VPC interface endpoints can be accessed through both intra-Region and inter-Region VPC peering connections.

VPC Gateway Endpoint connections can't be extended out of a VPC. Resources on the other side of a VPC peering connection in your VPC can't use the gateway endpoint to communicate with resources in the gateway endpoint service.

Reference: <https://docs.aws.amazon.com/systems-manager/latest/userguide/setup-create-vpc.html>

#### QUESTION 16

A company wants to monitor its AWS costs for financial review. The cloud operations team is designing an architecture in the AWS Organizations management account to query AWS Cost and Usage Reports for all member accounts.

The team must run this query once a month and provide a detailed analysis of the bill.

Which solution is the MOST scalable and cost-effective way to meet these requirements?

- A. Enable Cost and Usage Reports in the management account.  
Deliver reports to Amazon Kinesis.  
Use Amazon EMR for analysis.
- B. Enable Cost and Usage Reports in the management account.  
Deliver the reports to Amazon S3.  
Use Amazon Athena for analysis.
- C. Enable Cost and Usage Reports for member accounts.  
Deliver the reports to Amazon S3.  
Use Amazon Redshift for analysis.
- D. Enable Cost and Usage Reports for member accounts.  
Deliver the reports to Amazon Kinesis.  
Use Amazon QuickSight for analysis.

**Answer: C**

**Explanation:**

<https://docs.aws.amazon.com/cur/latest/userguide/what-is-cur.html>

If you are an administrator of an AWS Organizations management account and do not want any of the member accounts in your Organization to set-up a CUR you can do one of the following: (Recommended) If you've opted into Organizations with all features enabled, you can apply a Service Control Policy (SCP). Note that SCPs only apply to member accounts and if you want to restrict any IAM users associated with the management account from setting up a CUR, you'll need to adjust their specific IAM permissions. SCPs also are not retroactive, so they will not deactivate any CURs a member account may have set-up prior to the SCP being applied. Submit a customer support case to block access to billing data in the Billing console for member accounts. This is a list of organizations where the payer account prevents member accounts in its organization from viewing billing data on the Bills and Invoices pages. This also prevents those accounts from setting up Cost and Usage Reports. This option is only available for organizations without all features enabled. Please note that if you have already opted into this to prevent member accounts from viewing bills and invoices in the Billing Console, you do not need to request this access again. Those same member accounts will also be prevented from setting up a Cost and Usage Report.

#### QUESTION 17

A company has an application that ingests incoming messages. These messages are then quickly consumed by dozens of other applications and microservices. The number of messages varies drastically and sometimes spikes as high as 100,000 each second. The company wants to decouple the solution and increase scalability.

Which solution meets these requirements?

- A. Persist the messages to Amazon Kinesis Data Analytics.  
All the applications will read and process the messages.
- B. Deploy the application on Amazon EC2 instances in an Auto Scaling group, which scales the number of EC2 instances based on CPU metrics.
- C. Write the messages to Amazon Kinesis Data Streams with a single shard.  
All applications will read from the stream and process the messages.
- D. Publish the messages to an Amazon Simple Notification Service (Amazon SNS) topic with one or more Amazon Simple Queue Service (Amazon SQS) subscriptions.  
All applications then process the messages from the queues.

**Answer: D**

**Explanation:**

"SNS Standard Topic"

Maximum throughput: Standard topics support a nearly unlimited number of messages per second.

<https://aws.amazon.com/sns/features/>

"SQS Standard Queue"

Unlimited Throughput: Standard queues support a nearly unlimited number of transactions per second (TPS) per API action.

<https://aws.amazon.com/sqs/features/>

/

### QUESTION 18

A company that operates a web application on premises is preparing to launch a newer version of the application on AWS. The company needs to route requests to either the AWS-hosted or the on-premises-hosted application based on the URL query string. The on-premises application is not available from the internet, and a VPN connection is established between Amazon VPC and the company's data center. The company wants to use an Application Load Balancer (ALB) for this launch.

Which solution meets these requirements?

- A. Use two ALBs: one for on-premises and one for the AWS resource.  
Add hosts to each target group of each ALB.  
Route with Amazon Route 53 based on the URL query string.
- B. Use two ALBs: one for on-premises and one for the AWS resource.  
Add hosts to the target group of each ALB.  
Create a software router on an EC2 instance based on the URL query string.
- C. Use one ALB with two target groups: one for the AWS resource and one for on premises.  
Add hosts to each target group of the ALB.  
Configure listener rules based on the URL query string.
- D. Use one ALB with two AWS Auto Scaling groups: one for the AWS resource and one for on premises.  
Add hosts to each Auto Scaling group.  
Route with Amazon Route 53 based on the URL query string.

**Answer: C**

**Explanation:**

<https://aws.amazon.com/blogs/aws/new-advanced-request-routing-for-aws-application-load-balancers/>

The host-based routing feature allows you to write rules that use the Host header to route traffic to the desired target group.

Today we are extending and generalizing this feature, giving you the ability to write rules (and route traffic) based on standard and custom HTTP headers and methods, the query string, and the source IP address.

## Thank You for Trying Our Product

### Passleader 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.passleader.com/all-products.html>



Microsoft



ORACLE



CITRIX



JUNIPER  
NETWORKS



EMC<sup>2</sup>  
where information lives

**10% Discount Coupon Code: ASTR14**