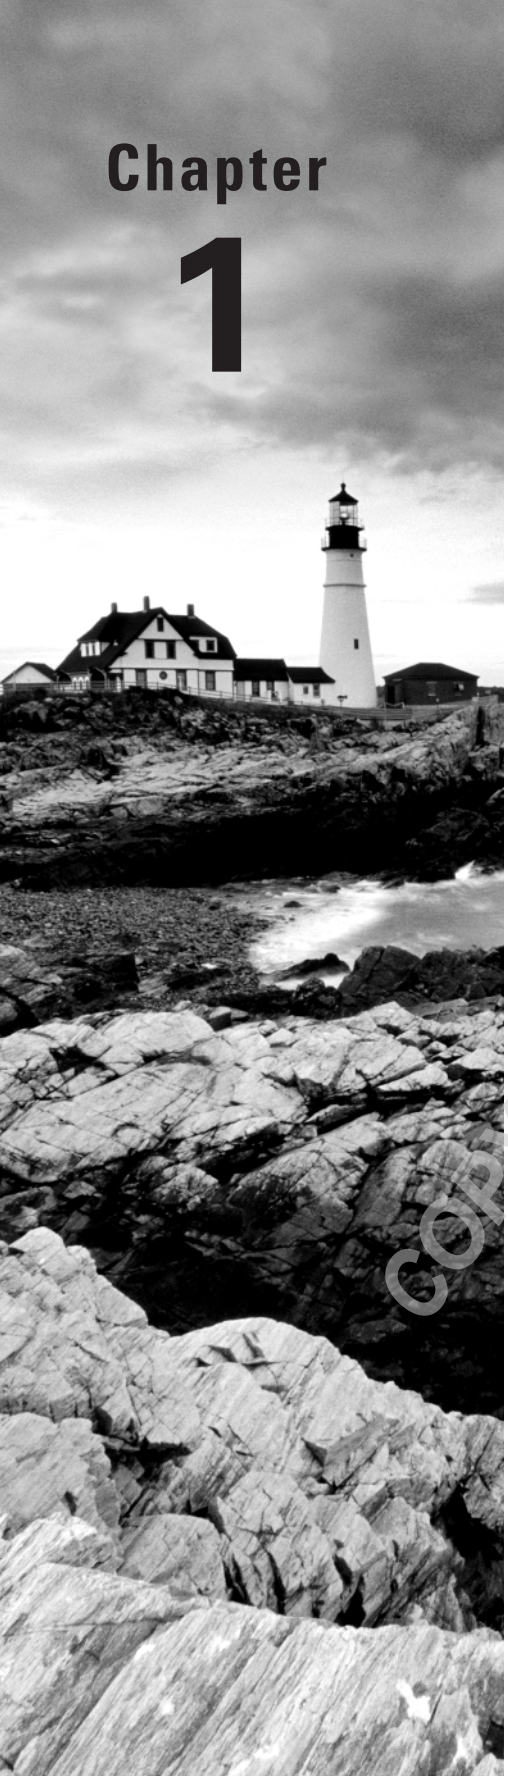

Chapter

1

Introduction to Systems Operations on AWS

THE AWS CERTIFIED SYSOPS ADMINISTRATOR - ASSOCIATE EXAM TOPICS COVERED IN THIS CHAPTER MAY INCLUDE, BUT ARE NOT LIMITED TO, THE FOLLOWING:

Domain 4.0: Deployment and Provisioning

✓ **4.1 Demonstrate ability to build the environment to conform with the architectural design**

✓ **4.2 Demonstrate ability to provision cloud resources and manage implementation automation**

Content may include the following:

- How to deploy cloud services
- Familiarity with three-tier architectures
- Deploying serverless architectures

Domain 6.0: Security

✓ **6.1 Ensure data integrity and access controls when using the AWS platform**

Content may include the following:

- AWS shared responsibility model
- AWS Cloudtrail
- Amazon EC2 Security Groups
- Network access control lists (ACLs)

Domain 7.0: Networking

✓ **7.1 Demonstrate the ability to implement networking features on AWS**

Content may include the following:

- Amazon Virual Private Cloud (Amazon VPC)

Systems Operators

You are a *systems operator*, and it is your job to keep your application environments running at maximum performance at all times. Just as a pit crew enables the racecar driver to win a race, systems operators are the pit crew—they help end users function successfully in their day-to-day jobs. You are an AWS systems operator, and this book will help you obtain the AWS Certified SysOps Administrator - Associate certification.

Deploying Systems

You might find yourself manually installing common, off-the-shelf packages on standalone instances. You might be coordinating an enterprise-wide effort to embrace fully-automated continuous deployment/continuous integration. Wherever you are on that spectrum, the responsibility to get it running in the first place falls on your shoulders.

However, deployment comprises much more than initializing systems. As enterprises evolve from monolithic application servers to container services, micro services, and serverless architectures, keeping up with the continuous stream of service updates requires attention and automation that you must manage.

Monitoring Systems

You might have a wall of monitors, all rendering real-time data on the environments in your care. You might have fully-automated alert functions that respond to changes in behavior, repairing or replacing failing parts and keeping you informed of these adjustments.

Nonetheless, you are monitoring much more than just network latency or CPU consumption. You have analytic engines that trace patterns in user behaviors—both consumers and employees. Your bots constantly review log files, looking for unusual activity and notifying you of anomalies.

Optimizing Systems

As a systems operator, you are your company's best agent for maximizing performance because your analytics help you choose the correct infrastructure configuration, the optimal storage methods, and the best possible customer outcome.

By 123net - Own work, CC BY-SA 3.0, https://commons.wikimedia.org/w/index.php?curid=17384917

However, you do more than optimize for speed; you optimize for cost. By using elastic environments, your environment not only automatically scales out during peak demand to minimize latency, but it also automatically scales in later to minimize spend. You manage an environment that is highly utilized every hour of every day.

Fortifying Systems

Things break and systems go offline, but you don't let that keep you up at night. You maintain highly available architectures: systems that detect failed components and automatically switch over, replacing and restoring as needed without interruption of service to your consumers.

But your availability methods cover more than single regions and multiple Availability Zones. Systems operations on AWS involves using multi-region and hybrid methods when needed to ensure continuity of operations no matter what Mother Nature throws at you.

Securing Systems

The combination of security groups, access control lists, and private networks in concert with native tools such as Amazon CloudFront and AWS Shield, help your environment stand up to the most sinister of attacks.

Threats don't always come from the outside, however. You know that the most dangerous vector is the internal attack. That's why you have meticulously employed a policy of compartmentalized, restricted privilege sets so that no one can step into unauthorized territory, along with detailed Application Programming Interface (API) logging that reports on all actions to provide comprehensive control over your assets.

AWS Certified SysOps Administrator - Associate

As detailed in the introduction to this chapter, AWS systems operators focus on a wide range of responsibilities. The *AWS Certified SysOps Administrator - Associate certification* is engineered to test your knowledge of systems operations domains. This book not only explains the domains on the exam, but it walks you through the different aspects of AWS with which you must be familiar in order to be successful as an AWS systems operator.

The test is organized into seven domains of relatively equal weight:

1. Monitoring and Metrics
2. High Availability
3. Analysis
4. Deployment and Provisioning
5. Data Management
6. Security
7. Networking

As you explore individual AWS architectures and services, it is important to note that many of the AWS products have operational considerations that apply to most, if not all, seven domains.

Which AWS Services Should You Study?

The simple answer is, "all of them."

AWS is constantly evolving and adding new offerings. As of this writing, AWS has more than 90 unique services. Each one has security, data, monitoring, and availability considerations. As an AWS systems operator, you are tasked with understanding those considerations along with how to optimize the service for performance and cost. The next few chapters in this book walk you through the service categories, explain how those services are addressed from an operational perspective, and discuss what you should study.

With more than 90 services and approximately 55 questions, mathematically not every service can be addressed in the certification exam. Commonly used services might appear in many different questions, although services with more specific use cases are much less likely to appear.

For example, when studying the storage products, you must understand the options found in Amazon Simple Storage Service (Amazon S3), Amazon Elastic Block Store (Amazon EBS), and Amazon Glacier. You can expect these services to appear in questions throughout all of the domains. In contrast, AWS Snowmobile could be on the test, but because it is used only in a few specific scenarios, statistically it is unlikely to appear more than once, if at all.

The best rule of thumb is to look at common reference architectures. If you see services in those architectures, plan on them being integral to the test. However, do not discount other services; everything is fair game.

The following section provides specific reference architectures that you can use as you plan on what services to study.

Reference Architecture: The Three-Tier Design

One of the earliest cloud-native architectures used is the three-tier design, which includes the following:

- A front-end web server layer
- An application middle layer
- A database layer

In many cases, the first two layers might be fronted, or decoupled, with elastic load balancers.

Introduction to the Three-Tier Design

The model of a three-tier architecture was introduced in the late 1990s. It was an evolution from a two-tier architecture (client/server), which was an evolution from a monolithic (mainframe-based) architecture. One of the original drivers for a three-tier architecture was the desire to implement a web-based interface to existing applications, which were currently being accessed via a command-line interface (CLI).

The focus of this model is on application architecture. Each application has its own unique architecture, which exists independently of any other application.

Web Tier

The *Web Tier* is the front end to the application. It accepts the request from the user and passes that request to the Application Tier. It takes the response from the Application

Tier and presents it back to the user. The format of the response is controlled at this tier, whether it is an HTML document, a CSV file, a PDF file, or some other format.

This tier has no direct access to the Database Tier, and it should be decoupled from any processes happening in the Application Tier or the Database Tier.

Application Tier

The *Application Tier* is a middleware tier where the internal business logic resides. It responds to requests from the Web Tier and communicates directly with the Database Tier. The Application Tier operates and scales independently of the other tiers.

Database Tier

The *Database Tier* is a back-end tier where the databases manage the state of the application. This tier should only be accessed by the Application Tier. It processes requests from the Application Tier and provides responses back to the Application Tier.

Sample Scenario

To better prepare you for the exam, this book references a few sample architectures. These are provided to give a framework to the discussions. Although the problem we might be addressing is specific, the services we use are universal to most architectures on AWS.

Real World Scenario

Three-Tier Architecture

The Challenge

An application runs an order management system for a global company. The application will manage inventory, customer records, and orders in an integrated system.

Some of the system requirements include flexibility to adjust to changing needs. It must be scalable to handle variable customer loads. It must have separate front-end and processing layers to allow User Interface (UI) development to be isolated from business logic programming.

It must be cost effective. In addition to scalable web and application instances, it should leverage native, cost-effective services such as elastic load balancing and Amazon S3.

The environment must be secure. Steps should be taken to ensure that all traffic is properly protected in transit and at rest. All access must be controlled and monitored at all times. All critical data must be stored in durable, highly-available systems, protected against node failure.

The Solution

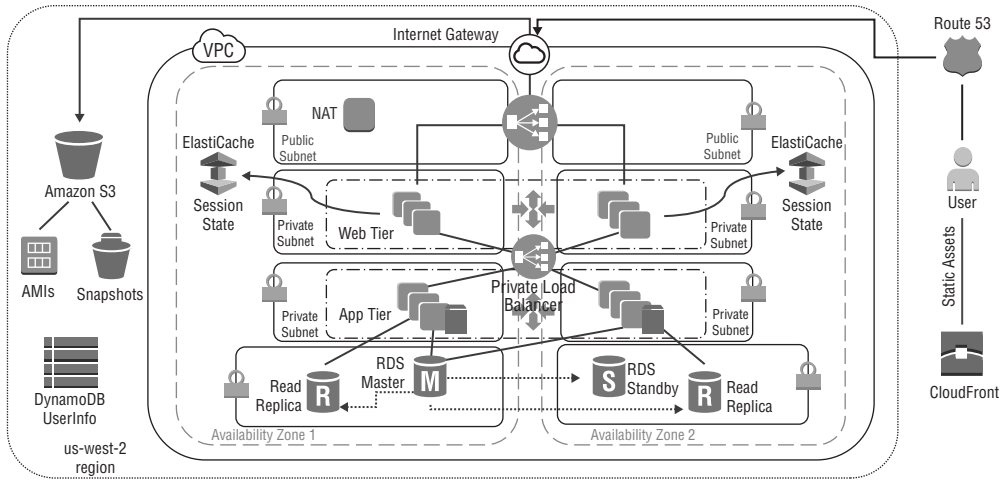

As we examine the pieces of the solution, we start by breaking down the components of the architecture. Then we focus on how systems operators interact with the individual pieces and begin thinking about how those pieces fit into the certification exam.

Environment

Architectures live inside *AWS Regions*; in this scenario, in us-west-2 (Oregon, United States). Regions are made up of multiple *Availability Zones*, which provide the foundation for highly available architectures. Although this is a systems operation exam, it is critical to understand the nature of AWS Regions and Availability Zones.

Each AWS Region is a separate geographic area. Each AWS Region has multiple, isolated locations known as *Availability Zones*. *AWS Regions* and *Availability Zones* are discussed in Chapter 5, "Networking."

Networking

Networking components start inside the AWS Region with Amazon Virtual Private Cloud (Amazon VPC). *Amazon VPC* is a private network in the AWS Region that isolates all traffic from the millions of other applications running in AWS. A deep dive into Amazon VPC (and the rest of its components) is found in Chapter 5.

Amazon VPC is divided into *subnets*; all assets running in your Amazon VPC are assigned to a subnet. Unlike on-premises subnetting decisions that can affect latency between servers, Amazon VPC subnets only affect access. Access between subnets is

controlled through *network Access Control Lists (nACLs)*, and access in and out of Amazon VPC is controlled through attached gateways. In this scenario, the only gateway is the *Internet Gateway (IGW)*, and it allows traffic to and from external (public IP) sources.

By granting route table access to the gateway only to specific subnets, ingress and egress can be tightly controlled. In this scenario, public subnets indicate IGW access. Without IGW access, the subnets become private; that is, they are accessible only to private IP networks.

To learn about the other gateways that could be leveraged to create hybrid or other private architectures, refer to Chapter 5.

Security groups are often part of the networking discussion. They provide stateful firewalls that operate at the Hypervisor levels for all individual*Amazon Elastic Compute Cloud (Amazon EC2)* instances and other Amazon VPC objects. In this scenario, we potentially have seven different security groups:

Public Elastic Load Balancing The only security group that allows full public access

Web Tier Amazon EC2 This accepts traffic only from public Elastic Load Balancing.

Private Elastic Load Balancing This accepts traffic only from Web Tier Amazon EC2.

Application Tier Amazon EC2 This accepts traffic only from private Elastic Load Balancing.

Amazon ElastiCache This accepts traffic only from Application Tier Amazon EC2.

Amazon Relational Database Service (Amazon RDS) This accepts traffic only from Application Tier Amazon EC2.

Network Address Translation (NAT) This is used only for internally initiated outbound traffic.

By specifically stacking security groups in this manner, you can provide layers of network security that surround the database portion of the three-tier design.

Compute

In this scenario, you use traditional compute methods, such as Linux servers running on Amazon EC2. Amazon EC2 comes in many sizes (how many CPUs, how much memory, how much network capacity, and so on), known as *instances*. Based on the Amazon Machine Image (AMI), each Amazon EC2 instance can run a wide range of Linux- or Windows-based operating systems as well as preinstalled software packages. Amazon EC2 instances also support runtime configuration as required.

The requirements for the scenario include scalable solutions. AWS provides Auto Scaling as an engine that can take predefined launch configurations and dynamically add or remove instances from the web or the Application Tier based on metrics.

Details on Amazon EC2, Auto Scaling, and other compute resources are found in Chapter 4, "Compute."

Database

Amazon RDS runs in your Amazon VPC on Amazon EC2. You select the database engine and version (MySQL, Oracle, Postgres, and so forth) and the configuration (the size of the Amazon EC2 instance, which subnets to use, how often to take backups, and so on). Amazon RDS takes care of the infrastructure of the instances and the engine; your database administrator (DBA) takes care of the database schema and data.

This scenario also includes *Amazon DynamoDB*, a native NoSQL engine optimized for consistent low latency, high availability, and strongly consistent reads and writes. Unlike Amazon RDS (or do-it-yourself databases running on Amazon EC2), Amazon DynamoDB operates at the regional level through API access only.

For details on how Amazon DynamoDB and other databases function, refer to Chapter 7, "Databases."

Storage

This scenario looks at storage in three different areas: the block storage used by the Amazon EC2 instances, the object storage keeping all of the media as well as backups and AMIs, and the caching storage used by Amazon CloudFront.

Amazon EBS is durable, persistent block storage used by most Amazon EC2 and Amazon RDS instances. It provides drive space for boot volumes and data volumes. Additionally, AWS provides ephemeral storage for many Amazon EC2 instance types through instance storage. Deciding which one to use becomes an operational value judgment, one that compares speed, persistence, and cost.

Object storage is provided by Amazon S3. *Amazon S3*, like Amazon DynamoDB, operates at the regional level outside Amazon VPC. It is only accessed through API commands that your operations team controls with fine-grained precision. Highly cost-effective and massively durable, Amazon S3 provides web-enabled storage for content as well as protected storage for database backups and AMI storage.

Amazon CloudFront is the *AWS content delivery network service (CDN)*. This application leverages Amazon CloudFront to cache content close to consumers in order to improve performance (reduce latency) and reduce costs.

Storage systems, including shared file systems, the Amazon Elastic File System (Amazon EFS), and cold storage via Amazon Glacier, are discussed in Chapter 6, "Storage."

User Management

Although not drawn in the sample three-tier architecture diagram, user management becomes one of the critical elements of the AWS operational design. Operator access is controlled through *AWS Identity and Access Management (IAM)*. IAM maintains control over validating authentication methods (passwords, access keys, and so on) and then grants access to authenticated operators.

Because everything in AWS is accessed through APIs, IAM becomes a comprehensive tool for controlling all permissions to AWS services and resources.

For established enterprise customers, IAM can be integrated with existing directory systems via AWS Directory Service.

AWS IAM controls access to AWS services and resources. It does not control access to the Amazon EC2 operating system or application-level authentication. For more details, refer to the shared responsibility model in Chapter 3, "Security and AWS Identity and Access Management (IAM)."

Security, Monitoring, and Deployment

Security is integral to every part of the AWS platform. This means that security is part of each piece of the architecture.

There are some specific AWS security tools, such as Amazon Inspector, Amazon VPC Flow Logs, Amazon CloudWatch Logs, and others which provide a more focused toolset that the AWS operations team can leverage to ensure the security profile of the AWS application. These and many other tools are discussed in Chapter 3.

Monitoring of critical systems is provided by *Amazon CloudWatch*, which provides visibility into metrics that happen on the Customer side of the shared responsibility model. Thousands of metrics across more than 90 services keep track of everything from CPU consumption to latency, queue depths, and so on.

AWS CloudTrail records every API call in the AWS system, including:

- Who made the API call
- When the API call was performed
- Where the API call originated
- The result of the API call

These records and other log files are processed through Amazon CloudWatch Logs, which analyze text data for patterns that trigger alerts and corresponding actions.

Automated deployment methods ensure that human error does not disrupt rollouts or updates to production or sandbox environments. *AWS CloudFormation* turns infrastructure plans into code, allowing your operations team to build and tear down entire systems in a single action. Refer to Chapter 8, "Application Deployment and Management," for more details.

Key Products: Three-Tier Design

As described above, the three-tier architecture consists of a web front end, an application layer, and database layer. In addition to the compute, storage, and database resources, additional AWS infrastructure may need to be deployed. Refer to Table 1.1 for a list of key products.

TABLE 1.1 Key Products: Three-Tier Architecture

Tools to Enable Hybrid Cloud Architectures	Description
AWS Regions and Availability Zones	Amazon EC2 is hosted in multiple locations world-wide. These locations are composed of regions and Availability Zones. Each region is a separate geographic area. Amazon EC2 provides you with the ability to place resources, such as instances and data, in multiple locations.
Availability Zones	Within a region are two or more Avalability Zones. The Availability Zone is isolated, but the Availability Zones in a region are connected through low-latency links.
Edge Locations	Edge Locations = AWS Lambda@Edge Amazon CloudFront, Amazon Route 53, AWS Shield, and AWS WAF services that are offered at AWS Edge Locations.
Hybrid cloud architecture	Integration of on-premises resources with cloud resources
Amazon Route 53	Amazon Route 53 is a highly available and scalable cloud Domain Name System (DNS) web service.
Amazon CloudFront	Amazon CloudFront is a web service that speeds up distribution of your static and dynamic web content to your users. CloudFront delivers your content through a worldwide network of data centers called *edge locations*.
Amazon VPC	A virtual private cloud (VPC) is a virtual network dedicated to your AWS account. It is logically isolated from other virtual networks in the AWS cloud. You can launch your AWS resources, such as Amazon EC2 instances, into your VPC.
Internet Gateways	An Internet gateway is a horizontally-scaled, redundant, and highly available VPC component that allows communication between instances in your VPC and the Internet.
Subnets	A subnetwork or subnet is a logical subdivision of an IP network.
Route tables	A route table is a set of rules that is used to determine where data packets traveling over an Internet Protocol (IP) network will be directed.

TABLE 1.1 Key Products: Three-Tier Architecture *(continued)*

Tools to Enable Hybrid Cloud Architectures	Description
Amazon EC2 Security groups	A security group acts as a virtual firewall that controls the traffic for one or more instances.
AWS Elastic Load Balanacing	Elastic Load Balancing automatically distributes your incoming application traffic across multiple targets, such as Amazon EC2 instances. It monitors the health of registered targets and routes traffic only to the healthy targets. Elastic Load Balancing supports two types of load balancers: Application Load Balancers and Classic Load Balancers.
Amazon Elastic Compute Cloud (Amazon EC2)	Amazon Elastic Compute Cloud (Amazon EC2) is a web service that provides secure, resizable compute capacity in the cloud.
Auto Scaling	Auto Scaling helps you maintain application availability and allows you to scale your Amazon EC2 capacity up or down automatically according to conditions you define.
Amazon Relational Database Service (Amazon RDS)	Amazon Relational Database Service (Amazon RDS) is a web service that makes it easier to set up, operate, and scale a relational database in the cloud.
Amazon DynamoDB	Amazon DynamoDB is a fully-managed NoSQL database service that provides fast and predictable performance with seamless scalability.
Amazon ElastiCache	Amazon ElastiCache is a web service that makes it easy to deploy, operate, and scale an in-memory data store or cache in the cloud.
Amazon Simple Storage Service (Amazon S3)	Amazon S3 is object storage with a simple web service interface to store and retrieve any amount of data from anywhere on the web.
Amazon Elastic Block Store (Amazon EBS)	Amazon EBS provides persistent block storage volumes for use with Amazon EC2 instances in the AWS Cloud.
Amazon Elastic File System (Amazon EFS)	Amazon EFS is a file storage service for use with Amazon EC2. Amazon EFS provides a file system interface, file system access semantics (such as strong consistency and file locking), and concurrently-accessible storage for up to thousands of Amazon EC2 instances.
Amazon Glacier	Amazon Glacier is an extremely low-cost storage service that provides secure, durable, and flexible storage for data backup and archival.

Tools to Enable Hybrid Cloud Architectures	Description
AWS Identity and Access Management (IAM)	AWS IAM is a web service that helps you securely control access to AWS resources for your users. You use IAM to control who can use your AWS resources (authentication) and what resources they can use and in what ways (authorization).
Active Directory Connector	AD Connector is designed to give you an easy way to establish a trusted relationship between your Active Directory and AWS.
Web identity federation	AWS IAM supports identity federation for delegated access to the AWS Management Console or AWS APIs. With identity federation, external identities (federated users) are granted secure access to resources in your AWS account without having to create IAM users.
Amazon CloudWatch	Amazon CloudWatch is a monitoring service for AWS cloud resources and the applications you run on AWS. You can use Amazon CloudWatch to collect and track metrics, collect and monitor log files, set alarms, and automatically react to changes in your AWS resources.
Amazon CloudWatch Logs	You can use Amazon CloudWatch Logs to monitor, store, and access your log files from Amazon Elastic Compute Cloud (Amazon EC2) instances, AWS CloudTrail, and other sources. You can then retrieve the associated log data from CloudWatch Logs.
Amazon VPC Flow Logs	VPC Flow Logs is a feature that enables you to capture information about the IP traffic going to and from network interfaces in your VPC. Flow log data is stored using Amazon CloudWatch Logs.
Amazon Inspector	Amazon Inspector is an automated security assessment service that helps improve the security and compliance of applications deployed on AWS. Amazon Inspector automatically assesses applications for vulnerabilities or deviations from best practices.
Amazon S3 Access Logs	In order to track requests for access to your bucket, you can enable access logging. Each access log record provides details about a single access request, such as the requester, bucket name, request time, request action, response status, and error code, if any.
AWS CloudTrail	AWS CloudTrail is a web service that records API calls made on your account and delivers log files to your Amazon S3 bucket.
AWS CloudFormation	AWS CloudFormation is a service that gives developers and businesses an easy way to create a collection of related AWS resources and provision them in an orderly and predictable fashion.

TABLE 1.1 Key Products: Three-Tier Architecture *(continued)*

Tools to Enable Hybrid Cloud Architectures	Description
AWS Elastic Beanstalk	AWS Elastic Beanstalk makes it even easier for developers to deploy and manage applications in the AWS Cloud quickly. Developers simply upload their application, and Elastic Beanstalk automatically handles the deployment details of capacity provisioning, load balancing, Auto Scaling, and application health monitoring.
AWS OpsWorks Stacks	AWS OpsWorks Stacks lets you manage applications and servers on AWS and on-premises. With OpsWorks Stacks, you can model your application as a stack containing different layers, such as load balancing, database, and application server.

It may seem like a daunting list, but this represents the core services (the toolset) that all AWS systems operators need to understand fully. As with any craft, it is important to use the right tool for the right job. You could use a torque wrench to smooth wet concrete, but of course there are much more appropriate tools for that task. Knowing the wide variety of AWS tools available to you is just as important.

Reference Architecture: The Serverless Design

As application design continues to evolve, individual instances are replaced with *container services*. Container services eventually are replaced by the final abstraction: *serverless architectures*.

There are many variations of serverless architectures. Rather than assume a generic use case, let's look at a specific scenario that might be used by your operations team.

Real World Scenario

Serverless Architectures

The Challenge

In this scenario, we want to find a better way to track the number of outstanding security updates on our production fleet. A serverless solution would be ideal, because we would not be adding any servers to maintain and we would only be paying for the compute time of the AWS Lambda functions.

The Solution

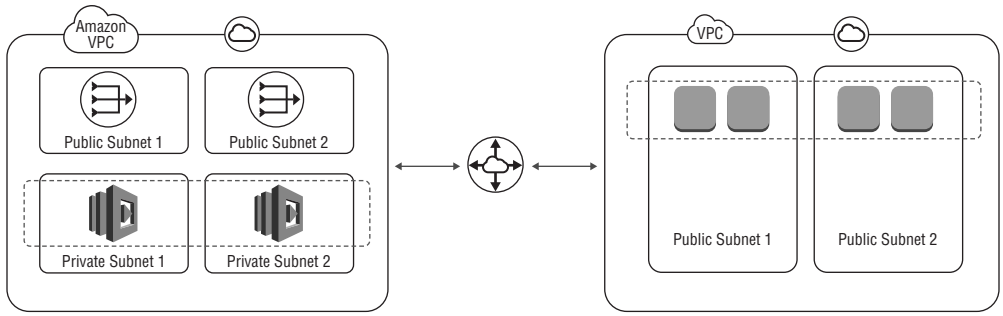

Python code executing in AWS Lambda on a regular schedule will use the Secure Shell (SSH) protocol to query for outstanding security updates on production instances. Python code (running anywhere) can use the AWS Boto Software Development Kit (SDK) to query Amazon EC2 for a list of specially tagged instances. The Python code establishes an SSH connection to the instances, and it executes a small script to find the number of required security updates. After you have this information, you can present it to the systems operations team as a tag on the instances, again using the AWS Boto SDK.

Networking

The *AWS Lambda functions* run in their own Amazon VPC. We establish Amazon VPC peering between the two Amazon VPCs to allow network connections between the AWS Lambda function and the production Amazon EC2 instances. This requires the creation of routing tables to direct the traffic between the two Amazon VPCs.

Security and Authentication

The AWS Lambda function must authenticate at two different levels: when the function queries the Amazon EC2 APIs via the Boto SDK and when the function establishes an SSH connection to the operating system on the production instances. AWS Lambda functions are configured with an IAM role and policy, which grants access to query the Amazon EC2 APIs. SSH authentication uses a Rivest-Shamir-Adleman (RSA) public/private key authentication. The AWS Lambda function has the private portion on the key. The Linux operating system on the production instances is configured with the public portion of the key. The operating system uses the public key to authenticate the SSH connection being initiated from the AWS Lambda function (see Figure 1.1).

FIGURE 1.1 Lambda function interacting with the Amazon EC2 API and EC2 instances

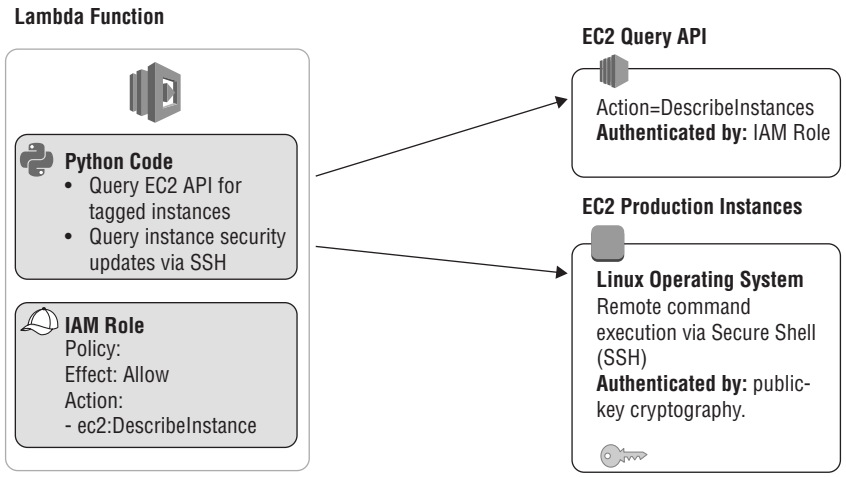

Lambda supports these runtime versions: Node.js, Java, and .Net Core. For more information, see Chapter 4.

Let's take an extra step to secure the private portion of the SSH key. This key is used by the AWS Lambda function to prove that it is allowed to SSH into the production instances and execute a script—so it is very important to keep secrets secret! The secret key is encrypted using the *AWS Key Management Service (AWS KMS)* and stored in Amazon S3. For the AWS Lambda function to retrieve the key from Amazon S3 and decrypt with AWS KMS, you must update the IAM policy associated with the AWS Lambda function. More information on cryptography is provided in Chapter 3. (See Figure 1.2.)

FIGURE 1.2 AWS KMS operations with Lambda

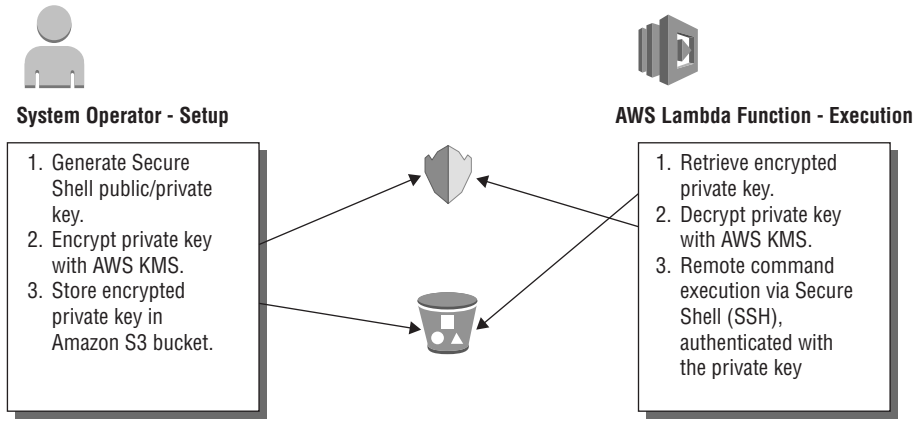

Who is allowed to access the encrypted private key in Amazon S3? Who is allowed to decrypt it? This is determined by the IAM policies in the AWS application.

Where and how do we apply network firewall type rules? The AWS Lambda function will be communicating to the production Amazon EC2 instances on the SSH port 22. Let's apply the least privilege principle here and ensure that only the AWS Lambda function is able to connect on port 22. We do this by creating security groups for both the production instances and the AWS Lambda function.

Key Product: Serverless Design

Many of the same services used in the three-tier architecture are used in the serverless design. Here are some of the unique services leveraged by this serverless architecture:

TABLE 1.2 Key Products: Serverless Design

AWS Product	Description
AWS Lambda	AWS Lambda lets you run code without provisioning or managing servers.
AWS Lambda@Edge	Lambda@Edge, now in Preview, allows you to write functions deployed to the AWS network of Edge locations in response to Amazon CloudFront.
AWS Key Management Service (AWS KMS)	AWS KMS is a managed service that makes it easy for you to create and control the encryption keys used to encrypt your data.
Amazon S3 web hosting	You can host a static website on Amazon S3. On a static website, individual web pages include static content. They may also contain client-side scripts. By contrast, a dynamic website relies on server-side processing, including server-side scripts such as PHP, JSP, or ASP.NET. Amazon S3 does not support server-side scripting.
Amazon API Gateway	Amazon API Gateway is a fully managed service that makes it easy for developers to publish, maintain, monitor, and secure APIs at any scale.
Amazon Kinesis	Amazon Kinesis is a platform for streaming data on AWS, offering powerful services to make it easy to load and analyze streaming data and also providing the ability for you to build custom streaming data applications for specialized needs.
Amazon Simple Queue Service (Amazon SQS)	Amazon Simple Queue Service (Amazon SQS) is a fully-managed message queuing service for reliably communicating among distributed software components and microservices—at any scale.

TABLE 1.2 Key Products: Serverless Design *(continued)*

AWS Product	Description
Amazon Simple Notification Service (Amazon SNS)	Amazon SNS is a web service that coordinates and manages the delivery or sending of messages to subscribing endpoints or clients.
Amazon Simple Email Service (Amazon SES)	Amazon SES is a low-cost solution for sending automated emails, such as order confirmations, shipping notices, order status updates, policy changes, password resets, and other messages that keep your customers informed.
AWS Web Application Firewall (AWS WAF)	AWS WAF gives you control over which traffic to allow or block to your web applications by defining customizable web security rules.
AWS Shield	AWS Shield is a managed Distributed Denial of Service (DDoS) protection service that safeguards web applications running on AWS.

Summary

Preparing for the certification exam requires comfort with a wide range of AWS services. One of the best ways to get comfortable is to use the AWS services themselves. Make sure that as part of your study, you take the time to create an account on AWS, log in to the AWS Management Console, launch the products, and get used to managing the various options. Practice builds the mental muscle memory that will give you the confidence in your answers.

Now that you know what types of architectures you will be dealing with and which products deserve the majority of your focus, let's start looking through the various service families covered throughout the AWS Certified SysOps Administrator - Associate exam.

Exam Essentials

Each chapter in this book ends with a list of important concepts to study. This list is not comprehensive, as the material is covered in the chapter itself, but the concepts are a good place to do a quick review of important testing areas. Every chapter ends with a useful tip from AWS trainers who specialize in helping people pass their certification exams. Look to these tips for good test-taking strategies that complement your core AWS knowledge.

Key Pieces to Study

Understand how AWS Regions and Availability Zones work to provide geographic distribution of services. Know how to deploy your environment across multiple Availability Zones and how to use Amazon CloudFront to take advantage of AWS edge locations.

Understand the shared responsibility model and that it is foundational to understanding how to secure your environment in AWS. Know which parts of any given service are managed by AWS and which parts you are responsible for securing.

Understand how the IAM engine separates the authentication layer from the authorization process. Be familiar with the way that credentials are presented to AWS when an API is called.

Test Taking Tips

Time management is key for this exam. You only have 80 minutes—don't waste them all on a question that has you stumped. Mark it for later review and move on. You will often be surprised that, when you come back to it later, the answer will be clear.

There is no penalty for wrong guesses. Make sure that you enter an answer for every question, even if you have no idea what the right answer might be. You won't pass the exam if you guess every question, but it never hurts to try on the few that you might not know.

The AWS Certified SysOps Administrator - Associate exam is not designed to give you trick questions. If one answer seems obviously right, but another answer might be correct under special circumstances, go with the obvious answer. Dr. Theodore Woodward's aphorism for his University of Maryland medical interns applies here: "If you hear hoof beats, think of horses not zebras."

Multiple-choice questions require all answers to be correct; there is no partial credit for getting a portion correct. Pay extra attention to those questions when doing your review.

Plan on leaving time at the end of the exam for review. Even if you think you know an answer, you can mark it and return to it when you are done with the exam. Go through each one of those marked questions to make sure that you are still confident with those answers. Just be careful not to overthink your answer (remember "horses not zebras").

Many questions have answer sets that are combinations of two pairs of answers. In AWS, everything is an API. In the next chapter, you will learn how to work with APIs and SDKs. So let's start our engines and get on with the nitty gritty of working with AWS Services!

Review Questions

1. Which AWS Cloud service allows you to gain system-wide visibility into resource utilization, application performance, and operational health?
   - **A.** Amazon CloudWatch
   - **B.** AWS OpsWorks
   - **C.** AWS Identity and Management (IAM)
   - **D.** AWS CloudTrail
2. Which AWS Cloud service enables you to capture information about the IP traffic going to and from network interfaces in your VPC?
   - **A.** Amazon CloudWatch
   - **B.** AWS OpsWorks
   - **C.** AWS CloudFormation
   - **D.** Amazon VPC Flow Logs
3. Which AWS Cloud service enables governance, compliance, operational auditing, and risk auditing of your AWS account?
   - **A.** Amazon CloudWatch
   - **B.** AWS CloudTrail
   - **C.** Amazon Simple Storage Service (Amazon S3) Access Logs
   - **D.** Amazon Elastic Compute Cloud (Amazon EC2) Security Groups
4. What is the term used for an environment that extends an existing on-premises infrastructure into the cloud to connect cloud resouces to internal systems?
   - **A.** Scatter architecture
   - **B.** Multi-location architecture
   - **C.** Hybrid cloud architecture
   - **D.** There isn't a term for this type of architecture.
5. Which of the following services acts as a virtual firewall that controls the traffic for one or more instances?
   - **A.** Network Access Control Lists (nACLs)
   - **B.** Security Groups
   - **C.** Availability Zones
   - **D.** Amazon Virtual Private Cloud (Amazon VPC)
6. A three-tier architecture is comprised of which of the following layers? (Choose three.)
   - **A.** Database layer
   - **B.** Front-end web server layer
   - **C.** Security layer
   - **D.** Application layer

7. Each AWS region is composed of two or more locations that provide you with the ability to introduce high availability, fault tolerance, and/or scale to your applications. What are these locations called?
   - **A.** Data centers
   - **B.** Edge locations
   - **C.** Compute centers
   - **D.** Availability Zones

8. What AWS Cloud service is designed to give you an easy way to establish a trusted relationship between your Active Directory and AWS?
   - **A.** Amazon Elastic Compute Cloud (Amazon EC2)
   - **B.** AWS Key Management Service (AWS KMS)
   - **C.** Amazon Virtual Private Cloud (Amazon VPC)
   - **D.** Active Directory Connector

9. What AWS Cloud service provides a logically isolated section of the AWS Cloud where systems operators can launch AWS resources into a virtual network they defined?
   - **A.** Amazon Virtual Private Cloud (Amazon VPC)
   - **B.** Amazon Route 53
   - **C.** Availability Zones
   - **D.** Security Groups

10. You manage a fleet of web servers hosted on Amazon Elastic Compute Cloud (Amazon EC2). Most, if not all, of the websites are static in nature. What AWS Cloud service can host a static website, thus replacing servers?
    - **A.** Amazon Elastic Compute Cloud (Amazon EC2)
    - **B.** Amazon Simple Storage Service (Amazon S3)
    - **C.** Amazon Route 53
    - **D.** Amazon API Gateway