PassQueen
a.k.a dumpsqueen.com

EC-Council Certified Cloud Security Engineer (CCSE)

ECCouncil 312-40

Version Demo

Total Demo Questions: 16

Total Premium Questions: 167

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Cloud Computing and Security Concepts	15
Topic 2, Cloud Computing Threats, Risks, and Countermeasures	2
Topic 3, Data Security in Cloud	15
Topic 4, Application Security in Cloud	7
Topic 5, Cloud Infrastructure Security	39
Topic 6, Cloud Security Operations	61
Topic 7, Cloud Security Governance, Compliance, and Risk Management	28
Total	167

QUESTION NO: 1

Dave Allen works as a cloud security engineer in an IT company located in Baltimore, Maryland. His organization uses cloud-based services; it also uses the Network Watcher regional service to monitor and diagnose problems at the network level. It contains network diagnostic and visualization tools that help in understanding, diagnosing, and obtaining visibility into the network in a cloud environment. This service helped Dave in detecting network vulnerabilities, monitoring network performance, and ensuring secure cloud operations. Which of the following cloud service providers offers the Network Watcher service?

A. Google

B. Azure

C. IBM

D. AWS

ANSWER: B

Explanation:

Azure is correct because Network Watcher is a Microsoft Azure regional service for monitoring, diagnosing, and gaining operational visibility into Azure network resources. It provides a collection of network diagnostic and visualization capabilities that cloud security and operations teams can use to investigate connectivity issues, observe traffic flows, validate routing and network security configuration, and monitor network behavior. For example, Network Watcher includes tools such as connection troubleshooting, IP flow verification, next-hop analysis, packet capture, topology views, and Network Security Group flow logs. These functions help an engineer identify conditions that can affect workload availability, performance, and security posture within an Azure virtual network environment. The service is enabled and used in the context of an Azure region, aligning with the question's description of a regional network-monitoring service. Microsoft's documentation identifies Network Watcher specifically as an Azure service and describes its monitoring and diagnostic toolset for Azure infrastructure. See the Azure Network Watcher overview and How Azure Network Watcher works for the supported capabilities and regional-service model.

QUESTION NO: 2

An AWS customer was targeted with a series of HTTPS DDoS attacks, believed to be the largest layer 7 DDoS reported to date. Starting around 10 AM ET on March 1, 2023, more than 15,500 requests per second (rps) began targeting the AWS customer ' s load balancer. After 10 min, the number of requests increased to 2,50,000 rps.

This attack resembled receiving the entire daily traffic in only 10s. An AWS service was used to sense and mitigate this DDoS attack as well as prevent bad bots and application vulnerabilities. Identify which of the following AWS services can accomplish this.

A. AWS Amazon Direct Connect

B. Amazon CloudFront

C. AWS Shield Standard

D. AWS EBS

E. AWS WAF

ANSWER: E

Explanation:

AWS WAF is the service that fits all of the requirements in the scenario. It protects web applications and APIs by inspecting HTTP/S requests before they reach protected resources such as an Application Load Balancer, Amazon CloudFront distribution, or API Gateway API. AWS WAF can identify and block unusually high request rates through rate-based rules, making it suitable for mitigating application-layer DDoS activity such as the reported HTTPS request flood. AWS documented the March 2023 event as a 2.5 million requests-per-second Layer 7 DDoS attack mitigated using AWS WAF protections.

Beyond volumetric request filtering, AWS WAF provides managed rule groups that address known application vulnerabilities and common web exploits. Its AWS Managed Rules offerings also include AWS WAF Bot Control, which classifies and controls bot traffic, allowing organizations to block, challenge, monitor, or otherwise manage unwanted automated requests. These capabilities directly satisfy the combined need for Layer 7 DDoS mitigation, bad-bot controls, and application-vulnerability protection. AWS Shield is an important DDoS-protection service, but the scenario's explicit bot and web-application vulnerability requirements identify AWS WAF as the intended service. See the AWS Security Blog report on the 2.5 million rps attack and the AWS WAF Bot Control documentation.

QUESTION NO: 3

Alex Hales works as a cloud security specialist in an IT company. He wants to make his organization's business faster and more efficient by implementing Security Assertion Mark-up Language (SAML) that will enable employees to securely access multiple cations with a single set of credentials. What is SAML?

A. It is a YAML-based authentication and authorization standard

B. It is an HTML based authentication and authorization standard

C. It is a XML based authentication and authorization standard

D. It is a JSON based authentication and authorization standard

ANSWER: C

Explanation:

It is a XML based authentication and authorization standard. Security Assertion Markup Language (SAML) is an open standard that uses XML to represent and exchange security assertions between an identity provider and a service provider. These assertions convey information such as a user's authenticated identity, relevant attributes, and authorization-related data. In a typical enterprise single sign-on workflow, an employee first authenticates with the organization's identity provider. The identity provider then creates a signed SAML assertion that the target cloud application, acting as the service provider, can validate and use to establish the employee's session without requiring a separate application-specific password. This federated approach centralizes identity management, improves the user experience, and supports controlled access to multiple SaaS or cloud services using one set of organizational credentials. XML is fundamental to the SAML specification: SAML protocol messages, assertions, bindings, and metadata are defined in XML, and XML Signature is commonly used to protect assertion integrity and establish trust between participating entities. The OASIS SAML technical overview describes SAML as an XML-based framework for communicating authentication, attribute, and authorization information. See the OASIS SAML standard and the SAML 2.0 Core specification.

QUESTION NO: 4

Assume you work for an IT company that collects user behavior data from an e-commerce web application. This data includes the user interactions with the applications, such as purchases, searches, saved items, etc. Capture this data, transform it into zip files, and load these massive volumes of zip files received from an application into Amazon S3. Which AWS service would you use to do this?

A. AWS Migration Hub

B. AWS Database Migration Service

C. AWS Kinesis Data Firehose

D. AWS Snowmobile

ANSWER: C

Explanation:

AWS Kinesis Data Firehose is the appropriate managed delivery service for continuously ingesting high-volume event data from an application and loading it into Amazon S3. A Firehose delivery stream accepts streaming records from producers,

buffers them according to configurable size or time thresholds, and writes the resulting objects to an S3 bucket without requiring the team to operate ingestion servers, manage batches, or build delivery retry logic. This maps well to e-commerce behavior events such as searches, saved items, and purchases, where data arrives continuously and must be collected at scale.

Firehose also supports processing records with an AWS Lambda transformation before delivery, enabling application events to be modified, enriched, or reformatted as required. For S3 destinations, it supports delivery-side compression, including GZIP, which reduces stored object size and transfer overhead. Thus, the service provides the managed capture, transformation, buffering, and S3-loading pipeline described in the question. AWS documents its Amazon S3 destination behavior at Amazon Data Firehose delivery to Amazon S3 and its transformation workflow at Transform source data in Amazon Data Firehose.

QUESTION NO: 5

Michael Keaton has been working as a cloud security specialist in a multinational company. His organization uses Google Cloud. Keaton has launched an application in nl-standard-1 (1 vCPU, 3.75 GB memory) instance.

Over the past three weeks, the instance has had low memory utilization. Which of the following machine type

switching is recommended for Keaton?

A. g1-small (0.5 vCPU, 1.7 GB memory)

B. n1-standard-2 (2 vCPU, 7.5 GB memory)

C. f1-micro (0.2 vCPU, 614 MB memory)

D. n1-standard-1 (1 vCPU, 3.75 GB memory)

ANSWER: A

Explanation:

g1-small (0.5 vCPU, 1.7 GB memory) is the appropriate rightsizing choice when an application running on an N1 standard instance has consistently low memory utilization. The existing N1 standard configuration provides 3.75 GB of memory, so three weeks of low observed use indicates that a smaller memory allocation may satisfy the workload while reducing the cost of idle capacity. The g1-small shared-core machine type provides 1.7 GB of memory, making it a practical downsize where the application does not require dedicated CPU performance and its actual memory demand remains comfortably below that amount. Rightsizing should be based on sustained monitoring data rather than a brief utilization dip, and the workload should be observed after the change to confirm that memory pressure, latency, and error rates remain acceptable. Google Cloud describes shared-core machine types as economical options for small, non-resource-intensive workloads, which fits an application with persistently low memory use. Before resizing, the administrator should verify peak-memory behavior and any application-specific minimum-memory requirement. See Google Cloud's machine type and resource documentation and its guidance on changing an instance machine type.

QUESTION NO: 6

Ryan has worked as a senior cloud security engineer over the past five years in an IT company. His organization uses Google cloud-based services because it provides live migration of VM. improved performance, robust security, better pricing compared to competitors. Ryan is using Cloud Endpoints to protect and manage APIs. Using Cloud Endpoints, ho is controlling access to APIs and validating every call with web tokens and Google API keys. Which of the following web tokens can validate every call in Cloud Endpoints?

A. SAML

B. JSON Web Tokens (JWTs)

C. XML

D. HTML

ANSWER: B

Explanation:

JSON Web Tokens (JWTs) are the token format Cloud Endpoints uses to authenticate API callers. A JWT is a compact, URL-safe credential containing signed claims, such as the issuer, audience, expiration time, and user or service identity. When an API is configured with an authentication provider in its OpenAPI specification, Cloud Endpoints validates the JWT presented in the request's `Authorization: Bearer` header before forwarding the request to the backend service. Validation includes checking the token signature against the provider's published JSON Web Key Set, confirming that the token was issued by a trusted issuer, and ensuring its audience matches the API configuration. These checks allow Endpoints to make an authorization decision for every protected API invocation without the backend having to independently validate the caller's identity. Google API keys can additionally identify the calling project and enforce API-key restrictions, while JWTs provide the authenticated identity and signed claims needed for secure access control. Google's Cloud Endpoints documentation describes configuring JWT authentication providers and validating JWTs for OpenAPI-based services: Authenticate users with JWTs in Cloud Endpoints. The JWT standard is defined in RFC 7519.

QUESTION NO: 7

Teresa Palmer has been working as a cloud security engineer in a multinational company. Her organization contains a huge amount of data; if these data are transferred to AWS S3 through the internet, it will take weeks. Teresa's organization does not want to spend money on upgrading its internet to a high-speed internet connection. Therefore, Teresa has been sending large amounts of backup data (terabytes to petabytes) to AWS from on-premises using a physical device, which was provided by Amazon. The data in the physical device are imported and exported from and to AWS S3 buckets. This method of data transfer is cost-effective, secure, and faster than the internet for her organization. Based on the given information, which of the following AWS services is being used by Teresa?

A. AWS Elastic Beanstalk

B. AWS Storage Gateway Volumes

C. AWS Storage Gateway Tapes

D. AWS Snowball

ANSWER: D

Explanation:

AWS Snowball is the correct service because it is an AWS Snow Family device designed for offline bulk data transfer when moving data over an internet connection would be impractical, slow, or costly. An organization requests a Snowball device from AWS, connects it to its on-premises environment, copies data onto the encrypted device, and returns it to AWS. AWS then imports the data into the designated Amazon S3 bucket. This workflow is specifically suited to transfers ranging from terabytes to petabytes, matching the scenario described. Snowball devices use strong encryption to protect data during transit, and AWS provides an end-to-end chain-of-custody process. The service avoids the need to upgrade a company's WAN connection solely to perform a one-time or periodic large-scale migration or backup transfer. AWS identifies Snowball Edge as a physical device used to transport large volumes of data into and out of AWS, including transfers involving Amazon S3. For an organization with insufficient available bandwidth and a large backup-data volume, using this physical transfer mechanism provides a practical, secure, and cost-effective ingestion path. See the AWS Snowball Edge Developer Guide and the AWS Snow Family overview.

QUESTION NO: 8

Ewan McGregor works as a cloud security engineer in a multinational company that develops software and applications for eCommerce companies. Owing to the robust services provided by AWS for developing applications and software, his organization migrated to the AWS cloud in 2010. To test whether it is possible to escalate privileges to obtain AWS administrator account access, Ewan attempt to update the login profile with regular user accounts. Which of the following commands should Ewan try to update an existing login profile?

A. aws iam update-login-profile -- user-name < password > -- password < username >

B. aws iam update-login-profile --user-name --password

C. aws iam update-login-profile -- user-name < password > -- password < username >

D. aws iam update-login-profile -- password < password > -- user-name < username >

ANSWER: B

Explanation:

`aws iam update-login-profile --user-name <username> --password <password>` is the correct AWS CLI syntax for changing the console password in an existing IAM user login profile. An IAM login profile contains the password-related settings that enable an IAM user to sign in to the AWS Management Console. The `--user-name` parameter identifies the IAM user whose profile is being changed, while `--password` supplies that user's new console password. The caller must have permission to perform the `iam:UpdateLoginProfile` action for the target user; whether such permission can be used for privilege escalation depends on the IAM policies, permissions boundaries, service control policies, and target-user restrictions configured in the AWS environment. AWS also supports password policy enforcement, so the supplied password must meet the account password policy when one is configured. In AWS CLI commands, option names must be written as `--user-name` and `--password`, without a space between the double hyphen and the parameter name. The official command reference documents the required password parameter and the optional user-name parameter for this operation. See the AWS CLI update-login-profile reference and the AWS IAM User Guide.

QUESTION NO: 9

An organization wants to detect its hidden cloud infrastructure by auditing its cloud environment and resources such that it shuts down unused/unwanted workloads, saves money, minimizes security risks, and optimizes its cloud inventory. In this scenario, which standard is applicable for cloud security auditing that enables the management of customer data?

A. Cloud Security Alliance

B. ISO 27001 & 27002

C. SOC2

D. NIST SP800-53 rev 4

ANSWER: B

Explanation:

ISO 27001 & 27002 is correct because these complementary international standards provide a structured basis for governing and auditing information security controls that protect customer data in cloud environments. ISO/IEC 27001 defines the requirements for establishing, operating, monitoring, reviewing, maintaining, and continually improving an information security management system (ISMS). This management-system approach supports a repeatable audit process for identifying cloud assets, assigning ownership, evaluating risks, maintaining an accurate inventory, and ensuring that unnecessary workloads are addressed according to organizational policy.

ISO/IEC 27002 provides implementation guidance for information-security controls, including asset management, access control, logging and monitoring, supplier relationships, information classification, and secure disposal. Applied to cloud services, these controls help an organization account for where customer data and workloads reside, determine whether their use is authorized, and manage the security and lifecycle of those resources. Together, the standards enable evidence-based cloud security auditing while connecting technical findings—such as unknown or unused infrastructure—to formal risk treatment and data-protection governance. ISO describes ISO/IEC 27001 as the ISMS requirements standard and ISO/IEC 27002 as guidance for information-security controls. See ISO/IEC 27001 information security management and ISO/IEC 27002:2022.

QUESTION NO: 10

Jerry Mulligan is employed by an IT company as a cloud security engineer. In 2014, his organization migrated all applications and data from on-premises to a cloud environment. Jerry would like to perform penetration testing to evaluate the security across virtual machines, installed apps, and OSes in the cloud environment, including conducting various

security assessment steps against risks specific to the cloud that could expose them to serious threats. Which of the following cloud computing service models does not allow cloud penetration testing (CPEN) to Jerry?

A. DBaaS

B. IaaS

C. PaaS

D. SaaS

ANSWER: D

Explanation:

SaaS is correct because this service model delivers a complete, provider-operated application to the customer. The provider retains operational control of the underlying infrastructure, virtualization layer, operating systems, and usually the application platform itself. Consequently, a customer does not have the administrative visibility or authority needed to perform the described penetration-testing activities against virtual machines, installed applications, and operating systems. Attempting active testing of those provider-managed components without explicit authorization could affect a shared service and fall outside the customer's permitted scope.

In a SaaS arrangement, the customer's security responsibility is principally focused on data, identities, access configuration, and safe use of the service. Security testing must therefore be limited to authorized customer-controlled areas and conducted under the SaaS provider's documented testing, vulnerability-disclosure, and acceptable-use processes. Any broader assessment requires written approval and a clearly defined scope from the provider. This reflects the cloud shared-responsibility model: the provider secures and operates the SaaS application and its supporting cloud stack, while the customer secures its data and access within that service. See the Microsoft shared responsibility model and AWS shared responsibility model.

QUESTION NO: 11

Cosmic IT Services wants to migrate to cloud computing. Before migrating to the cloud, the organization must set business goals for cloud computing as per the guidelines of a standard IT governance body. Which standard IT governance body can help the organization to set business goals and objectives for cloud computing by offering the IT governance named COBIT (Control Objective for Information and Related Technology)?

A. International Standards Organization (ISO)

B. Cloud Security Alliance (CSA)

C. Information System Audit and Control Association (ISACA)

D. Committee of Sponsoring Organizations (COSO)

ANSWER: C

Explanation:

Information System Audit and Control Association (ISACA) is correct because ISACA develops, maintains, and publishes COBIT, now commonly expanded as *Control Objectives for Information and Related Technologies*. COBIT is an IT governance and management framework that helps enterprises connect information and technology activities to stakeholder needs and strategic business objectives. For a cloud migration, this supports defining intended business outcomes, establishing governance structures, assigning accountability, managing risk, measuring performance, and ensuring that cloud-related investments deliver value. COBIT's goals cascade provides a structured way to translate enterprise goals into alignment goals and then into specific governance and management objectives. This makes it particularly useful before and during cloud adoption, when the organization must ensure that technology decisions, controls, risk tolerance, compliance needs, and service-management practices remain aligned with business priorities. ISACA provides COBIT resources and guidance for organizations seeking to govern digital and IT-enabled services, including cloud environments. The official ISACA COBIT overview is available at https://www.isaca.org/resources/cobit. Further information about ISACA's role in advancing digital trust, governance, audit, risk, and security is available at https://www.isaca.org/.

QUESTION NO: 12

Rachel McAdams works as a cloud security engineer in an MN

C.

A DRaaS company has provided a disasterrecovery site to her organization. The disaster recovery sites have partially redundant equipment with daily or weekly data synchronization provision; failover occurs within hours or days with minimum data loss. Based on this information, which of the following disaster recovery sites is provided by the DRaaS company to Rachel's organization?

A. Warm Site

B. Cold Site

C. Remote site

D. Hot Site

ANSWER: A

Explanation:

Warm Site is correct because it provides a partially prepared recovery environment that includes some preinstalled and operational infrastructure, while requiring additional configuration, restoration, or activation before production workloads can run after a disruption. The stated daily or weekly synchronization schedule aligns with a recovery design that accepts some data loss between synchronization events. Likewise, a failover time measured in hours or days is characteristic of a warm environment: the organization has equipment and a recovery location available, but the environment is not continuously running as an immediately usable production replica.

In disaster-recovery-as-a-service arrangements, a warm recovery environment balances resilience against operating cost. It is more capable than an empty facility because key systems and equipment are available, yet it avoids the expense of maintaining a fully active duplicate environment at all times. Recovery personnel may need to restore the most recent synchronized data, start services, complete configuration tasks, and validate the recovered applications before business operations resume. AWS describes warm standby as a scaled-down but functional version of a production environment that can be expanded during recovery; this reflects the same general recovery-site model: AWS Disaster Recovery: Warm Standby. For broader recovery planning guidance, see Microsoft Azure disaster recovery guidance.

QUESTION NO: 13

Cindy Williams works as a cloud security engineer in an IT company located in Seattle, Washington. Owing to the cost-effective security, governance, and storage features provided by AWS, her organization adopted AWS cloud-based services. Cindy would like to detect any unusual activity in her organization's AWS account. She would like to obtain the event history of her organization's AWS account activity for security analysis and resource change tracking. Which of the following AWS service enables operational auditing, compliance, governance, and risk auditing for her organization's AWS account?

A. AWS CloudFormation

B. AWS Security Hub

C. AWS Config

D. AWS CloudTrail

ANSWER: D

Explanation:

AWS CloudTrail is the AWS service designed to provide an event history of activity in an AWS account. It records events generated by actions taken through the AWS Management Console, AWS CLI, SDKs, and AWS service APIs. These records identify details such as the acting identity, action requested, source IP address, time, affected resources, and result of the action. This visibility gives security teams the audit trail needed to investigate unusual activity, analyze security events, and track changes to AWS resources.

CloudTrail directly supports operational auditing, security analysis, governance, risk auditing, and compliance requirements because organizations can review recent management events in Event history and configure trails or event data stores for ongoing retention, querying, and analysis. For example, CloudTrail logs can be delivered to Amazon S3 and integrated with monitoring and detection workflows, enabling an organization to retain evidence and investigate activity over time. AWS documents CloudTrail specifically as a service for governance, compliance, operational auditing, and risk auditing of an AWS account. See the AWS CloudTrail User Guide and the CloudTrail Event history documentation.

QUESTION NO: 14

A new public web application is deployed on AWS that will run behind an Application Load Balancer (ALB). An AWS security expert needs to encrypt the newly deployed application at the edge with an SSL/TLS certificate issued by an external certificate authority. In addition, he needs to ensure the rotation of the certificate yearly before it expires. Which of the following AWS services can be used to accomplish this?

A. AWS Snowball

B. AWS Certificate Manager

C. AWS Cloud HSM

D. Amazon Elastic Load Balancer

ANSWER: B

Explanation:

AWS Certificate Manager is the appropriate service because it integrates directly with an Application Load Balancer to provide HTTPS/TLS termination at the load balancer listener. ACM supports importing certificates that were issued by an external certificate authority, including the certificate, private key, and applicable certificate chain. Once imported, the certificate can be selected for the ALB's secure listener, allowing clients to establish encrypted connections to the public application endpoint.

For certificates obtained from an external CA, the organization remains responsible for obtaining the renewed certificate from that CA before expiration. The renewed certificate can then be imported into ACM, including through supported ACM APIs as part of a scheduled certificate-rotation process. ACM also provides certificate-expiration monitoring through Amazon EventBridge events and Amazon CloudWatch metrics, enabling operational automation and alerting before the yearly renewal deadline. This makes ACM the AWS service that centrally stores, deploys, and supports the lifecycle process for an externally issued TLS certificate used by an ALB. See the AWS documentation for importing certificates into ACM and ACM certificate renewal documentation.

QUESTION NO: 15

Maria Howell has been working as a senior cloud security engineer in an IoT manufacturing company. Her organization designs, develops, and tests IoT devices. It uses Microsoft Azure cloud-based services. Maria had no knowledge of data science and the various ML and AI models used for data analysis, but she would like to analyze the time-series data generated from IoT devices to monitor and identify abnormalities. Which of the following is an AI-based Azure service that can help Maria in monitoring and identifying the abnormalities in time series data without requiring any knowledge of machine learning?

A. Application Insights

B. Azure Sentinel

C. Cloud App Security

D. Anomaly Detector

ANSWER: D

Explanation:

Anomaly Detector is the Azure AI service designed specifically to identify unusual patterns, spikes, dips, level changes, and other anomalies in time-series data. It is appropriate for IoT manufacturing telemetry because device measurements, sensor readings, operational metrics, and environmental values are naturally captured as timestamped sequences. Maria can submit univariate or multivariate time-series data through the service API and receive anomaly results without having to create training pipelines, choose algorithms, engineer features, or understand the underlying machine-learning models.

The service provides prebuilt AI capabilities, including batch and real-time detection scenarios, so it supports both retrospective analysis of collected device data and ongoing monitoring of incoming telemetry. Results identify whether a data point is anomalous and can include contextual information that helps an application trigger alerts or investigation workflows. This managed approach matches the requirement for a cloud-based capability usable by someone without data-science or ML expertise. Microsoft describes the service and its time-series anomaly detection operations in the Azure AI Anomaly Detector overview and provides implementation guidance in the multivariate anomaly detection quickstart.

QUESTION NO: 16

A company is a third-party vendor for several organizations and provides them customized software and products to cater to their needs. It recently moved its infrastructure and applications on cloud. Its applications are not working on the cloud as expected. The developers and testers are experiencing significant difficulty in managing and deploying the code in the cloud. Which of the following will help them with automated integration, development, testing, and deployment in the cloud?

A. Vulnerability assessment tool

B. DevOps

C. SIEM

D. Dashboard

ANSWER: B

Explanation:

DevOps is the appropriate approach because it brings development and operations practices together to make software delivery repeatable, collaborative, and highly automated in cloud environments. In this scenario, the teams need a way to reliably manage code changes and move them through integration, testing, and deployment after migrating their applications and infrastructure. DevOps practices address this need through continuous integration, in which changes are frequently merged, built, and automatically tested, and continuous delivery or deployment, in which validated changes are released through automated pipelines. These pipelines can also provision cloud resources consistently through infrastructure as code, apply configuration controls, and provide monitoring feedback after release. This reduces manual deployment effort, improves consistency between environments, speeds delivery of fixes, and gives developers and testers a shared process for maintaining application quality. Cloud platforms commonly provide managed DevOps tooling for source control, build automation, testing, artifact management, release pipelines, and operational monitoring. Microsoft describes DevOps as the union of people, processes, and technology to continuously provide value, while its continuous integration guidance explains the automated build-and-test workflow central to this use case. See Microsoft: What is DevOps? and Microsoft: What is continuous integration?.