Google CDL Exam Glossary

Find more at BareMetalCyber.com

1. **Access control**
   Access control is the set of rules that decides **who** can do **what** with **which** cloud resources. On this exam it shows up as "pick the safest default," often separating authentication (proving identity) from authorization (granting permissions). Google Cloud

2. **Agility**
   Agility is an organization's ability to change direction quickly, ship updates faster, and respond to customer needs without long procurement cycles. The exam commonly frames agility as a business outcome of cloud adoption, and contrasts it with slower on-premises change cycles and rigid capacity planning. Google Cloud

3. **Application programming interface (API)**
   An API is a defined way for one system to request data or actions from another system, using a consistent contract. It matters because many "modernization" questions use API s as the bridge that lets teams reuse or extend legacy capabilities without rewriting everything at once.

4. **Application modernization**
   Application modernization is updating how an app is built, deployed, or operated so it fits modern delivery patterns, like managed services or containers. On the exam, the decision point is usually choosing modernization to improve speed, reliability, or maintainability rather than simply moving the same design to a new location.

5. **Artificial intelligence (AI)**
   AI refers to systems that perform tasks associated with human intelligence, such as recognizing patterns, understanding language, or making predictions. The exam tends to test when A I is appropriate (and when it is not) and how it connects to data quality, governance, and business value.

6. **Availability**
   Availability is the likelihood that a service is usable when needed, often discussed as uptime and resilience. On the exam, availability choices often hinge on selecting the right architecture idea (redundancy, regional design) rather than a single "bigger server" mindset.

7. **Autoscaling**
   Autoscaling is automatically adjusting capacity up or down based on demand so performance stays stable and cost stays controlled. It commonly appears as a cloud benefit tied to elasticity, and the trap is confusing it with manual provisioning or assuming scale automatically fixes poor design.

8. **BigQuery**
   BigQuery is Google Cloud’s fully managed analytics data warehouse used for analyzing large datasets with SQL-style queries. On the exam, it often represents the “analytics warehouse” choice compared to operational databases, and the key is matching it to reporting and insight workloads rather than transaction processing.

9. **Bigtable**
   Bigtable is a managed NoSQL wide-column database designed for very large scale, low-latency workloads. The exam usually uses it to test whether you can distinguish “operational, high-throughput, low-latency” storage needs from “warehouse analytics” needs.

10. **Billing account**
    A billing account is the container that owns payment responsibility and links costs to the right organization, projects, and teams. It matters because many cost-governance questions test whether you can separate technical usage from financial ownership and apply controls like budgets and chargeback thinking.

11. **Budget**
    A budget is a cost guardrail that sets an expected spending limit for a scope like a project or billing account. On the exam it shows up as a governance control, often paired with alerts, so teams notice unusual spend early instead of discovering it after the invoice.

12. **Business continuity**
    Business continuity is the ability to keep critical business functions running during and after disruptions. The exam typically tests continuity as a planning and design outcome, where the right answer emphasizes resilience patterns and clear recovery expectations, not just “back up everything.”

13. **Cloud computing**
    Cloud computing is delivering compute, storage, and services over the internet with on-demand access and managed operations. It matters because many questions

ask you to connect cloud characteristics to business outcomes like speed, scalability, and standardized security controls.

14. **Cloud deployment model**
    A cloud deployment model describes where workloads run and who owns the environment, such as public cloud, private cloud, hybrid, or multi-cloud. The exam often tests your ability to choose a model based on constraints like compliance, latency, data residency, or existing investments.

15. **Cloud service model**
    A cloud service model describes how much is managed for you, commonly framed as Infrastructure as a Service, Platform as a Service, and Software as a Service. It shows up as "who is responsible for what," where the key is understanding how responsibility shifts as services become more managed.

16. **Cloud Storage**
    Cloud Storage is Google Cloud's object storage for files and unstructured data, designed for durability and flexible access patterns. The exam uses it as the default "store objects" option, and a common confusion is mixing it up with block storage or databases meant for structured queries.

17. **Compliance**
    Compliance is meeting external and internal requirements, such as laws, regulations, and contractual obligations. On the exam, compliance is usually a decision filter that changes architecture, data handling, and access control choices, and it is often confused with security even though they are not identical.

18. **Container**
    A container packages an application with its dependencies so it runs consistently across environments. The exam tests containers as a modernization tool that improves portability and deployment consistency, but it also expects you to recognize they do not remove the need for security and governance.

19. **Cost optimization**
    Cost optimization is the practice of reducing waste and aligning spending to actual value delivered. Exam questions often present it as a set of choices like rightsizing, turning off unused resources, and using managed services, while avoiding the trap of sacrificing reliability to save money.

20. **Customer-managed encryption keys**
Customer-managed encryption keys are encryption keys controlled by the customer rather than fully controlled by the cloud provider. This term matters because key ownership and control often appears as a compliance or risk requirement, and the exam tests when stronger key control is appropriate versus unnecessary overhead.

21. **Data governance**
Data governance is the set of policies and controls that define how data is created, classified, accessed, retained, and used. On the exam it shows up as the “how do we stay in control” layer behind analytics and A I, and it is commonly confused with tooling rather than decision rules and accountability.

22. **Data lake**
A data lake is a storage approach that keeps large amounts of raw or lightly processed data so different teams can analyze it later. The exam often tests when a lake makes sense for flexibility and varied data types, and when a warehouse is better for consistent reporting and curated datasets.

23. **Data pipeline**
A data pipeline is the end-to-end flow that collects, moves, transforms, and delivers data from sources to destinations. It matters because many scenarios test your ability to pick a design that supports reliable ingestion and repeatable processing, not one-off copying.

24. **Data privacy**
Data privacy is protecting personal or sensitive data by limiting collection, use, exposure, and retention based on rules and consent. On the exam, privacy shows up as a constraint that affects identity access, logging, and data handling choices, and it is easy to confuse with security because they overlap but have different goals.

25. **Data residency**
Data residency is the requirement that data be stored or processed in specific geographic locations. The exam uses residency to drive region-selection decisions and governance controls, and the trap is choosing a technically convenient region that violates legal or contractual requirements.

26. **Data warehouse**
A data warehouse is a curated, structured repository designed for analytics and reporting with consistent definitions. The exam typically tests the warehouse as the

"single source of truth" for business reporting and contrasts it with operational databases or data lakes.

27. **Disaster recovery**
    Disaster recovery is the plan and capability to restore systems and data after a significant disruption. The exam often checks whether you can match recovery objectives to business needs, recognizing that higher resilience usually costs more and requires clearer design choices.

28. **Elasticity**
    Elasticity is the ability to automatically scale resources up and down as demand changes. It shows up as a core cloud benefit, and the common confusion is mixing elasticity with performance tuning or assuming elasticity replaces good capacity planning decisions.

29. **Encryption**
    Encryption is converting data into a protected form so only authorized parties with the right key can read it. On the exam, encryption commonly appears in questions about protecting data at rest and in transit, and the decision point is choosing it as a baseline control rather than an optional add-on.

30. **Identity**
    Identity is a digital representation of a user, service, or system that can be authenticated and granted permissions. The exam frequently uses identity as the start of security reasoning, where the trap is jumping to network controls without first ensuring identity and permissions are correct.

31. **Identity and access management (I A M)**
    I A M is the discipline and set of controls that define who can access resources and what actions they can take. On the exam, I A M is a frequent "first best answer" area because it addresses least privilege and clear accountability, and it is often tested as distinct from network security.

32. **Infrastructure as a Service (I a a S)**
    I a a S is a cloud service model where you manage operating systems and applications while the provider manages underlying hardware and basic virtualization. It matters because exam questions often test responsibility boundaries, especially around patching, configuration, and security hardening.

33. **Key management**
Key management is the creation, storage, rotation, and control of cryptographic keys used for encryption and signing. On the exam, it appears as a governance and risk topic where the key confusion is assuming "encrypted" automatically means "keys are well controlled."

34. **Latency**
Latency is the time delay between a request and a response, often felt as slowness by users or systems. The exam uses latency to steer choices like region placement, network design, or caching, and the trap is focusing only on compute size while ignoring distance and network paths.

35. **Least privilege**
Least privilege means granting only the minimum permissions needed to perform a task, for the shortest necessary time. It shows up as a core security principle on the exam, and the common mistake is choosing broad roles for convenience instead of scoped access with traceable approvals.

36. **Logging**
Logging is recording events and actions so teams can audit activity, troubleshoot issues, and investigate incidents. The exam tests logging as an operational and security control, and a common confusion is collecting logs without ensuring retention, searchability, or ownership for review.

37. **Managed service**
A managed service is a cloud offering where the provider runs and maintains much of the underlying infrastructure and operational tasks. On the exam, managed services are often the best fit when the goal is speed, reliability, and reduced operational burden, but the trap is assuming you lose all responsibility.

38. **Migration**
Migration is moving applications, data, or workloads from one environment to another, such as from on-premises to cloud. The exam often tests migration as a means to an end, where the right answer aligns the approach to goals like risk reduction, faster delivery, or modernization readiness.

39. **Multi-region**
Multi-region refers to designing systems to run across more than one region to improve resilience and availability. It matters because exam scenarios may require

higher continuity, and the key tradeoff tested is increased complexity and cost versus reduced outage impact.

40. **Network security**
    Network security is protecting data and services by controlling network paths, segmentation, and exposure to threats. On the exam, it commonly appears alongside identity controls, and the confusion is treating networks as the only security layer instead of one layer in a broader control set.

41. **Observability**
    Observability is the ability to understand what a system is doing by examining signals like logs, metrics, and traces. On the exam it shows up as the difference between reacting to outages and proactively detecting issues, and it is often tested as a capability that supports reliability and security monitoring.

42. **On-premises**
    On-premises refers to infrastructure and systems owned and operated by an organization in its own facilities or dedicated space. The exam uses on-premises as a contrast point for cloud benefits and constraints, and questions often test when hybrid approaches exist because not everything can move at once.

43. **Platform as a Service (P a a S)**
    P a a S is a service model where the provider manages the runtime, scaling, and much of the platform layer while you focus on code and data. It matters on the exam because it changes responsibility boundaries and usually improves speed to deliver, but you still must handle secure configuration and access.

44. **Project**
    A project is a logical container used to organize cloud resources, access controls, and billing in a structured way. On the exam, projects are frequently part of governance questions where scoping, ownership, and separation of environments reduce risk and improve cost clarity.

45. **Reliability**
    Reliability is the ability of a system to perform correctly and consistently over time, including recovering from failures. The exam tests reliability as a design and operations outcome, often requiring you to balance resilience, cost, and complexity rather than chasing the highest possible availability.

46. **Risk management**
    Risk management is identifying risks, evaluating their likelihood and impact, and choosing responses such as reducing, transferring, accepting, or avoiding them. On the exam it shows up in governance and security decisions, and the trap is treating risk as purely technical instead of business-driven.

47. **Scalability**
    Scalability is the ability of a system to handle growth in users, data, or workload without breaking or degrading unacceptably. The exam distinguishes scalability from elasticity by emphasizing long-term growth planning, and it often appears in modernization and architecture choices.

48. **Shared responsibility model**
    The shared responsibility model describes how security and compliance duties are divided between the cloud provider and the customer. It is heavily tested as a reasoning tool, especially when choosing managed services, and the common confusion is assuming the provider handles everything once something runs in the cloud.

49. **Software as a Service (S a a S)**
    S a a S is a model where the provider delivers a complete application and manages most operational responsibilities. The exam tests S a a S as the fastest path to capability when customization needs are low, and it often uses it to check if you understand the smallest operational burden among service models.

50. **Zero trust**
    Zero trust is a security approach that assumes no implicit trust and requires verification for each access request based on identity, device, context, and policy. On the exam it appears as a modern security posture tied to least privilege and strong identity controls, and the key confusion is thinking it is a single product rather than a set of principles applied consistently.