

Q1. What is Azure OpenAI and what is its primary purpose in its cloud ecosystem?

A: Azure OpenAI is a managed cloud service that handles a specific infrastructure concern (compute, storage, messaging, AI). Using it shifts operational burden to the cloud provider while you focus on application logic.

Q2. How does IAM work with Azure OpenAI?

A: Access to Azure OpenAI is controlled via IAM roles and policies. Attach the minimum required permissions to a service account or role. Avoid long-lived access keys; use instance roles or Workload Identity Federation instead.

Q3. How do you monitor Azure OpenAI in production?

A: Azure OpenAI emits metrics to the cloud provider's monitoring service (CloudWatch, Cloud Monitoring, Azure Monitor). Set alerts on key metrics (error rate, latency, capacity). Use distributed tracing to correlate across services.

Q4. How do you scale Azure OpenAI to handle variable load?

A: Azure OpenAI supports auto-scaling policies based on CPU, queue depth, or custom metrics. Set minimum and maximum capacity. Horizontal scaling (more instances) is generally preferred over vertical for cloud workloads.

Q5. How do you secure Azure OpenAI in production?

A: Place Azure OpenAI in a private subnet with no public endpoint where possible. Use VPC security groups or firewall rules to allow only required traffic. Encrypt data at rest and in transit. Rotate credentials regularly.

Q6. How do you manage configuration and secrets for Azure OpenAI?

A: Store sensitive values in a managed secret store (AWS Secrets Manager, GCP Secret Manager, Azure Key Vault). Inject secrets as environment variables at runtime, never bake them into container images or source code.

Q7. What are the main cost drivers for Azure OpenAI and how do you optimize them?

A: Cost in Azure OpenAI typically comes from compute hours, data transfer, storage, and request counts. Right-size instances, use reserved capacity for steady-state workloads, and enable lifecycle policies to expire old data.

Q8. How do you implement high availability with Azure OpenAI?

A: Deploy Azure OpenAI across multiple availability zones. Use managed replicas or active-active configurations. Implement health checks and automatic failover. Test resilience regularly with chaos engineering or failover drills.

Q9. How does Azure OpenAI integrate with a CI/CD pipeline?

A: CI/CD pipelines use the cloud CLI or SDK to deploy updates to Azure OpenAI after tests pass. Infrastructure-as-code (Terraform, CDK) defines Azure OpenAI configuration as version-controlled code deployed via the pipeline.

Q10. What are common pitfalls when first adopting Azure OpenAI?

A: Common mistakes include misconfigured IAM policies (too permissive), no cost alerting, deploying only in one availability zone, and missing backups. Read the service SLA carefully to understand what the cloud provider guarantees.