

104 AWS EC2 interview questions to hire top engineers

Questions

1. What is AWS EC2 and why would someone use it?
2. Can you explain the different types of EC2 instances and when you might choose one over another?
3. How do you launch an EC2 instance and what are the key steps involved?
4. What is an AMI, and how is it used when creating EC2 instances?
5. How do you connect to an EC2 instance once it's running?
6. What are security groups in EC2, and how do they help protect your instances?
7. Explain the difference between public and private IP addresses in the context of EC2.
8. What is Elastic IP and when would you need to use one?
9. How do you manage the storage attached to your EC2 instances? What are EBS volumes?
10. What are EC2 key pairs and why are they important for security?
11. How can you monitor the performance of your EC2 instances?
12. What are some ways to automate tasks related to EC2 instances, like starting and stopping them?
13. What is the difference between stopping, hibernating, and terminating an EC2 instance?
14. How do you handle backups and recovery for your EC2 instances?
15. Explain how you can scale your application using EC2 instances. What is auto-scaling?
16. What are placement groups and when would you use them?
17. How does EC2 integrate with other AWS services?
18. What are some best practices for securing EC2 instances?
19. Describe a scenario where you would use a dedicated host instead of a shared EC2 instance.
20. How do you optimize costs when using EC2 instances?
21. What is the purpose of instance metadata? How do you access it?
22. Can you explain the different EC2 pricing models (On-Demand, Reserved, Spot)?
23. How do you update software on an EC2 instance?
24. What steps would you take to troubleshoot high CPU utilization on an EC2 instance?
25. How do you configure an EC2 instance to run a specific application automatically when it starts?
26. Imagine your EC2 instance is super slow. How do you figure out what's making it slow, like checking its heartbeat?
27. Let's say you have a bunch of EC2 instances doing the same thing. How do you make sure traffic is spread evenly among them, like sharing toys fairly?
28. If you want to run a task automatically on your EC2 instance when it starts up, like setting an alarm clock, how would you do it?
29. What are some ways you can make your EC2 instances super secure, like putting a lock on your toy chest?
30. If your application needs different resources at different times of the day how can you configure EC2 to scale up and down based on demand?
31. How can you ensure your application running on EC2 is highly available, and what are some of the key architectural considerations to implement?
32. Explain the difference between EBS volume types like gp3, io2, and st1 and how would you choose the right one for different workloads?
33. How do you monitor the health of your EC2 instances and respond automatically to failures, like an instance crashing?
34. Describe the process of creating a custom AMI and why might you need to customize an AMI?
35. How can you use EC2 instance metadata and user data to configure your instances dynamically at launch?
36. What are placement groups in EC2, and how do they help in optimizing network performance for certain applications?
37. Explain the various methods for managing and distributing SSH keys for EC2 instances in a secure manner.
38. How do you handle long-running processes on EC2 instances so that they are fault-tolerant and can survive instance failures?
39. Discuss the challenges of running stateful applications on EC2 instances, and how you would address them?
40. What are the advantages and disadvantages of using Spot Instances for your workload, and how can you mitigate the risks?
41. Describe how you would implement a blue/green deployment strategy using EC2 instances?
42. How can you use AWS Systems Manager (SSM) to manage and automate tasks on your EC2 instances?
43. Explain how you would configure network interfaces and routing tables to set up a multi-tier architecture on EC2?
44. How do you handle patching and updates on EC2 instances to maintain security and stability?
45. What are some best practices for optimizing the cost of running EC2 instances, and what tools can help?
46. How would you go about troubleshooting high CPU utilization on an EC2 instance running a web application?
47. Explain how to integrate EC2 instances with other AWS services like S3, RDS, and Lambda in a secure and efficient manner.
48. Discuss how you would design a disaster recovery plan for your EC2-based applications?
49. What strategies can you use to optimize disk I/O performance for EC2 instances, particularly for database workloads?
50. How can you leverage EC2 Auto Scaling groups with lifecycle hooks to perform custom actions before or after instances launch or terminate?
51. How would you optimize EC2 costs for an application that experiences peak usage only during specific hours of the day?
52. Describe a scenario where you would choose a Dedicated Host over a Dedicated Instance, and explain why.
53. Explain how you can automate the patching of EC2 instances while minimizing downtime.
54. How would you design a disaster recovery strategy for EC2 instances across multiple AWS regions?
55. What are the key considerations when migrating a large, complex on-premises application to EC2 instances?
56. Explain the differences between EC2 placement groups and when you would use each type.
57. How do you monitor the health and performance of EC2 instances in a production environment, and what metrics are most important?
58. Describe how you would implement a blue/green deployment strategy for EC2-based applications.
59. How can you ensure that sensitive data stored on EC2 instances is protected at rest and in transit?
60. Explain how you would integrate EC2 instances with other AWS services like S3, RDS, and Lambda.
61. What are the advantages and disadvantages of using EC2 Image Builder for creating custom AMIs?
62. How would you troubleshoot a situation where an EC2 instance is experiencing high CPU utilization or memory pressure?
63. Explain how you can use EC2 Systems Manager to manage and automate tasks on your EC2 instances.
64. Describe a scenario where you would use EC2 Auto Scaling with instance weighting, and explain how it works.
65. How would you design a highly available and scalable web application architecture using EC2 instances and load balancing?
66. Explain the different types of EBS volume encryption and how they impact performance.
67. How can you use AWS CloudTrail to audit and monitor changes made to your EC2 instances?
68. Describe how you would implement a multi-tier application architecture using EC2 instances in different VPC subnets.
69. How do you handle license management for software running on EC2 instances?
70. Explain how you can use EC2 Fleet to request a combination of instance types and Availability Zones.
71. How would you optimize network performance between EC2 instances in the same or different Availability Zones?
72. Explain how you would implement centralized logging for EC2 instances across multiple AWS accounts.
73. How can you use EC2 Instance Connect to securely connect to your EC2 instances without using SSH keys?
74. Describe how you would implement a CI/CD pipeline for deploying applications to EC2 instances.
75. How would you diagnose and resolve network connectivity issues between EC2 instances and other AWS services?
76. Explain how you can use AWS Config to track the configuration of your EC2 instances and ensure compliance with your organization's policies.
77. How do you choose the appropriate EC2 instance type for a given workload, considering factors like CPU, memory, storage, and network performance?
78. Explain a scenario where you'd choose EC2 Dedicated Hosts over Dedicated Instances and why it's the better choice in that situation.
79. How can you ensure that your EC2 instances are launched in a specific Availability Zone and what are the implications of failing to do so?
80. Describe the steps to migrate an on-premises virtual machine to an EC2 instance, highlighting potential challenges and mitigation strategies.
81. Explain how you would monitor the performance of your EC2 instances to identify and resolve bottlenecks.
82. What are the key considerations when selecting the appropriate EBS volume type for a specific EC2 workload?
83. Detail a disaster recovery strategy for EC2 instances running a critical application, including RTO and RPO considerations.
84. How can you automate the process of patching and updating your EC2 instances?
85. Discuss the security best practices for securing your EC2 instances, including network configuration, access control, and encryption.
86. Describe a scenario where you would use EC2 Auto Scaling in conjunction with a load balancer to ensure high availability and scalability.
87. Explain how you would optimize the cost of running your EC2 instances without sacrificing performance or availability.
88. How do you configure and manage multiple network interfaces on an EC2 instance?
89. Describe how you would troubleshoot a situation where an EC2 instance is experiencing high CPU utilization.
90. What are the steps involved in creating a custom Amazon Machine Image (AMI) and why would you do so?
91. Explain how you would use AWS Systems Manager to manage and configure your EC2 instances at scale.
92. How can you integrate EC2 instances with other AWS services, such as S3, DynamoDB, and Lambda?
93. Explain the process of setting up and managing a bastion host for secure access to your EC2 instances.
94. Describe how you would use CloudWatch to monitor the health and performance of your EC2 instances.
95. What are the key considerations when choosing between different EC2 instance types for a specific workload?
96. Explain how you would use CloudFormation to automate the deployment of EC2 instances and other AWS resources.
97. How can you use EC2 Instance Connect to securely connect to your EC2 instances without using SSH keys?
98. Describe the process of setting up and managing a VPC Peering connection between two VPCs with EC2 instances in each.
99. Explain how you would use AWS Config to track changes to your EC2 instance configurations and ensure compliance.
100. How can you use placement groups to influence the placement of your EC2 instances for performance or availability?
101. Describe a scenario where you would use EC2 Fleet to launch a diverse set of EC2 instance types and sizes.
102. Explain how you would use AWS Certificate Manager (ACM) to manage SSL/TLS certificates for your EC2 instances.