AWS Advanced Networking Specialty Master Cheat Sheet

Exam Details [link]

The Advanced Networking exam is 170 hours and successful completion is judged on the following.

Content Outline:

- **Domain 1**: Design and implement hybrid IT network architectures at scale, *23%*
- **Domain 2**: Design and implement AWS networks, *29%*
- **Domain 3**: Automate AWS tasks, *8%*
- **Domain 4**: Configure network integration with application services, *15%*
- **Domain 5**: Design and implement for security and compliance, *12%*
- **Domain 6**: Manage, optimize, and troubleshoot the network, *13%*

Example questions provided by AWS can be found here.

Exam Resources

- **A large amount of AWS documentation**, there is a mountain of pages that will explain the nuances of the available AWS services, I will link to many of the pages I indexed when studying for the exam in my notes [link].
- **AWS Whitepapers**: [link]
    - *AWS Security Best Practices* [link]
    - *Security at Scale: Governance in AWS* [link]
    - *Overview of AWS Security - Network Security* [link]
    - *Security at Scale: Logging in AWS* [link]
    - *AWS Best Practices for DDoS Resiliency* [link]
    - *An Introduction to High Performance Computing on AWS* [link]
    - *Amazon Virtual Private Cloud Network Connectivity Options* [link]
    - *Integrating AWS with Multiprotocol Label Switching* [link]
    - *Best Practices for Deploying Amazon WorkSpaces* [link]

Study Notes

Topics:

Elastic Network Interfaces [link]

- You can associate multiple IPs to each network interface [link], know the quantities for each instance type [link].
- Each ENI can have multiple IPv6 addresses if the VPC has an associated IPv6 CIDR block [link].
- You can move an ENI between subnets, but not availability zones [link].
- Attaching two ENIs to the same instance in the same subnet can cause networking issues. If necessary use a second IPv4 address on the primary network interface instead.
- Cross-account ENIs are possible but they need to be whitelisted by AWS, this can be a process and automation bottleneck.

Elastic IPs [link]

- To stop being charged for an elastic IP address:
  - Release it from the account.
  - Associate it to a running instance.
- You can reclaim an Elastic IP via the CLI [link].

Elastic Network Adapters [link]

- Original introduction [link].
- Elastic Network Adapter will support speeds of up to 20Gbps (different from ENI).
- ENA Support needs to be flagged when registering an AMI.

Enhanced Networking [link]

- Enhanced Networking for Linux [link] and Windows [link].
- Guide for AWS Network Bandwidth [link].
- 5Gbps is the maximum speed for EC2 to EC2 traffic outside of a placement group.
- Intel network interface: 10Gbps in placement group.
- Elastic Network Adapter: 20gbps [link].

VPC Gateways [link]

- Requirements:
  - Route tables referencing prefixes w/ public IPs.

    - DNS resolution.
    - Proxy if accessing externally from VPC with private VIF or VPN.
- You can't have an endpoint for a VPC to connect to a service in any region.
- AWS keeps an up to date list of IP address for gateway services:
    - this is accessible from `ip-ranges.json` [link],
    - changes to the IP ranges can be subscribed to with the SNS topic `AmazonIpSpaceChanged` [link].

VPC Interfaces [link]

- Accessible between VPCs w/ private IPs.
- Accessible over DX.
- Not accessible over VPN, VPC peering.

PrivateLink [link]

- Introduction [link].
- A type of interface VPC endpoint, includes use of Network Load Balancer.
- A PrivateLink interface consumer interface can establish connections but the producer end cannot.
- PrivateLink limitations [link]:
    - No inter-region access,
    - Only IPv4 with TCP,
    - Need to work with AWS to align AZs for interface endpoints.

VPCs [link]

- Being able to perform CIDR arithmetic for subnets should be second [link] nature [link].
- VPCs can be between `/16` and `/28` - using the `10.0.0.0/8`, `172.16.0.0/12` and `192.168.0.0/16` ranges.
- Out of the box VPC range `172.31.0.0/16`.
- You can recreate the default VPC these days [link].
- Max five IGW's in a region.
- 5 addresses reserved per subnet:
    - Network address `.0`,
    - VPC router `.1`,
    - Reserved for DNS (only used for VPC `.2` address),
    - Reserved for future use, `.3`,
    - Reserved for broadcast `.255` (unused) [link].
- NAT gateway needs Elastic IP [link].
- VPC peering is now allowed amongst VPCs in different regions link.

- Cloudhub will not allow VPC to VPC peering, Cloudhub is for connecting on-prem to VPC.
- Internet Gateway attached can be denied from Organizations [link].
- Go with an agent based approach for IPS/IDS, no promiscuous mode [link].
- AWS does not support broadcast, multicast between instances, it is only possible with a VPN overlay using GRE tunnels.

EC2 [link]

- if you are creating an EC2 mail server, alert AWS about the IP so they know it isn't spamming.
- If you *reboot* (not stop/start an instance) you retain the public address and stays on the same host commputer [link].
- If you stop and start an EC2 instance you lose your public IP address [link].

Security Groups/NACLs [link]

- Security group max rules - 50, max security groups per EC2 instance - 5 (250 possible).
- A Security group cannot be applied to an Egress Only Gateway [link] or NAT Gateway [link].
- You require two subnets per Database subnet group.
- in case of a DDOS remove the default NACL in a VPC to temporarily cut off traffic.
- Ephemeral ports *1024 – 65535* are required outbound for return traffic for web servers, check this on firewalls.

Direct Connect and Virtual Interfaces

- `AWS cloudwatch list-metrics --namespace "AWS/DX"`
- Virtual interface needed for each VPC. Connected to Direct Connect.
- Hosted virtual interface: use by account different from the account that owns the direct connect.
- Sub 1Gbps -> hosted connection.
- One LOA per connection per datacentre. LAGs count as one.
- 100 prefixes announced over private VIF.
- The DX Gateway will allow access to multiple regions.
- Only way to stop billing, delete the direct connect connection, removing VIF wont do it.
- Can't access S3 endpoint on private VIF.
- Multi-Exit Discriminators (MED): [link]
- Differentiated Services Code Point (DSCP): [link] For classifying/managing network traffic QoS for Layer 3.

- Forwarding Equivalence Class (FEC): [link].

DX Requirements [link]

- BGP.
- BGP MD5 auth.
- single mode fibre *1000BASE-LX (1310nm)* for 1 Gbps ethernet, *10GBASE-LR* for 10Gbps with 802.1q VLANs [link].
- Auto-negotiation for the port for direct connect needs to be disabled.
- You cannot change the port speed of an existing connection.
- limit on BGP advertised routes per route table? 100.
- Lowest bandwidth on DX partners is 50 mbps.

VPN [link]

- IPSec and Encapsulating Security Protocol [link].
  - IP protocol 50, port 500 UDP for IPSec [link].
- Provides:
  - Data encryption across the internet.
  - Protection of data in transit over the internet.
  - Peer ID authorisation between VPN gateway and customer gateway.
  - Data integrity across the internet.
- AWS Cloudwatch can be used to monitor a VPN, but not in the sense of the monitor that keeps an IPSec tunnel open [link].
- You need a monitoring tool for a VPN tunnel to remain up.
- AWS VPN does not support 128 bit AES, does support 4-byte ASN.
- Static VPN: 50 routes IPv4, 50 routes IPv6 max.
- Dynamic VPN w/ BGP: 100 routes max.
- To run VPN over DX, you need to have a public VIF to access the VPN endpoints.
- Use multiple customer gateways and dynamic routing for highly available VPN (you can use multiple customer gateways).
- Can’t use S3 endpoint with VPN, can use Public VIF + VPN.

Routing Scenarios

- A more specific route is the only way to make AWS prefer a VPN over a Direct Connect connection.
- 2 VPCs with VPN tunnels back to on prem. Both have same address. Address this with VRF in your router for each network, then delete a VPN, create a placeholder, recreate the final VPN and delete the placeholder (described in [link]).
- AS_PATH prepending assists in setting active/passive mode.

- A public VIF can be used to reach public services in regions other than the connected one. This is only for North America with nonlocal public ranges advertised over BGP [link].
- For 2 DXs where you need active/passive, how do you force? set AS_PATH prepending on passive connection, prepending makes path less preferred.
- higher MED makes path less preferred, less desirable way to reduce priority than AS_PATH prepending.
- 2 VPN connections established, how do you make one preferred? More specific route.
- With VPN, direct connect pairing, adventise less specific prefix on VPN connection to favour direct connect.
- Understand routing priority [link] and VPN routing priority [link].
- Cloudhub: Create multiple cust gateways with a public IP address.
- Create VPN connection between each customer gateway and a common virtual private gateway.
- Route summarization for BGP routes where there are >100 routes.
- A more specific CIDR route makes the path more preferred.
- Peering between three accounts, where the two spokes have overlapping ranges.

Link Aggregation Groups [link]

- Operate in Active/Active mode.
- Can be a part of another LAG. Must be on the same AWS device.
- Is a layer 2 connection.
- LAGs with 2 direct connect connections, have 1 BGP sessions per VIF.

Maximum Transmission Units (MTU) [link]

- Jumbo frames inside the VPC, not inter-VPC.
- To maximise instance speed for inter and intra VPC speed have two ENIs, one for external with MTU of 1,500 bytes and one for internal with MTU of 9,001 bytes.
- Instance type can determine max MTU.
- `tracepath` needs UDP.
- To check MTU between 2 hosts? `tracepath` (available from iputils).
- VPN connections and traffic sent over an Internet gateway are limited to 1500 MTU. If packets are over 1500 bytes, they are fragmented, or they are dropped if the Don't Fragment flag is set in the IP header [link].

OSI Layers [link]/IP Suite [link]

- IPSec operates at the network/internet layer.

- TCP/UDP are transport layer.
- DNS/DHCP operate at the application layer.
- FTP/SMTP is application layer.
- BGP is application layer.
- AWS Shield operates at network and transport layer.

Placement Groups [link]

- Placement group:
  - instances all need to be started at the same time,
  - can span VPCs,
  - deployment strategies: cluster, spread,
  - You cannot move an existing instance into one.

Route53 [link]

- **CNAME**:
  - Not free for queries,
  - points to DNS records hosted anywhere,
- **Alias** [link]:
  - No charge for alias queries to AWS resources,
  - Only for AWS resources
  - Only visible via console or API
- CNAME record needs to be updated in domain name registrar [link].
- Creating a hosted zone creates the following records: name server, SOA [link].
- To make sure some nameservers are used in Route 53, in the API create a Reusable Delegation Set.
- Route53 has health checks that can check the health of other health checks.
- Route53 doesn't do UDP health checks (but does do HTTPS).
- instances launched in VPC get public DNS names? `enableDnsHostnames`.
- **Shuffle sharding**: shard hosted zones over multiple instances to mitigate DDOS attacks [link].
- **Anycast striping**: multiple instances can respond to the same IP address, anycast IP addresses can route you to closest anycast server [link].
- You can't create a Reusable Delegation Set in a console [link].
- White label servers: name servers in Route53 will be the same as the domain name of your hosted zone.

Elastic Load Balancers [link]

- ELB needs a subnet size of `/27`.

- Connection draining time for an Elastic Load Balancer, min: 1 second max: 1 hour, default: 5 min.
- SSL negotiation needs Security policy, SSL protcols, SSL ciphers, server order preference [link].
- No WAF on ELB classic, only ALB.
- Terminating SSL on the load balancer [link]
- ALB and ELB classic supports HTTP X-forwarded-for [link].
- Sticky sessions, available for ELB Classic and ALB.
- ALB, NLB allow target groups.
- `x-forwarded-for` header needed to see client in access logs.
- ELB Classic:
    - Proxy protocol only for Classic Load Balancer, this cannot be enabled from console,
    - Can use TCP/443,
    - Uses alias record [link],
    - Configurable idle timeout [link],
    - Default interval for logs? 60 minutes.

AWS Workspaces [link]

- Needs minimum MTU of 1200.
- Workspaces: 2 private, 1 public subnet.
- AWS Workspaces has 2 network interfaces.
- Workspaces uses Active Directory, options: Microsoft AD/AD connector/Simple AD.
- AWS Workspaces can use on-prem RADIUS for Multi-Factor Authentication [link].

Active Directory [link]

- Simple AD does not work with microsoft products.
- Simple AD supports less than 5,000 users.
- AD connector will cause too much authentication traffic.

Cloudfront [link]

- HTTPS cannot be used between cloudfront/S3.
- Cloudfront always forwards requests to S3 by using the protocol that viewers used to subnet requests.
- Cloudfront Origin Protocol Policy should be match viewer to get http or https.
- WAF -> Cloudfront, ALB integrations.

- CloudFront origin access identity: - create a special cloudfront user, apply read only perms for identity to the bucket via policies.
- CloudFront expiration date for URLs? use signed URLS.
- CloudFront Origin Access Identity, add to S3 bucket policy.
- Security groups can be set to only be open to CloudFront IP ranges [link].
- Shared secret custom headers in origin calls from AWS CloudFront [link].

DDOS & Environment Protection [link]

- DDOS can be TCP or UDP.
- AWS Shield Advanced is available (and is pricey) [link].
- To create a baseline for WAF traffic initialise the WAF in monitor mode to begin with.

Networking Tools

Some useful tools for networking of Linux instances in AWS.

Checking network performance between two servers.

- iperf3: Speedtest tool for TCP/UDP.
- ec2-net-utils: This package optionally automate the configuration of your network interfaces. Available for Amazon Linux only. Usage `yum install ec2-net-utils`.
- mtr: Combines the functions of the traceroute and ping programs in one network diagnostic tool.
- hping3: A command-line oriented TCP/IP packet assembler/analyzer. Works with TCP/UDP and ICMP.
- tcpdump: A packet analyser, allows the user to display TCP/IP and other packets being transmitted or received over a network.

Enable jumbo frames:

```
  sudo ip link set dev eth0 mtu 9001
```

Enable on restart:

/etc/dhcp/dhclient-eth0.conf

```
interface "eth0" {
    supercede interface-mtu 1500;
}
```

- Get instance MTU:

```
ip link show eth0 | grep mtu
```

Results

In the end I sat the exam and passed! The types of questions I encountered included:

- Architecting solutions given constraints involving resiliency, cost, performance, security and compliance requirements.
- Deciding between the different types of network access and siloing including proxying, peering, transitive setups and endpoints and diagnosing traffic routing edge cases.
- Specifics of ELBs, I had a lot of questions involving the differences between the types of load balancer.
- Troubleshooting involving Direct Connect.
- A large portion of Route 53 questions, troubleshooting name resolution between AWS, on premises and the internet.
- Gateway and Interface logic including when you would use PrivateLink or cross account ENIs
- Gotchas surrounding multicast and packet sniffing, `ip-ranges.json`.
- Cloudfront features, integrations and networking edge cases.
- A lot of troubleshooting of firewall boundaries.
- High performance networking, placement groups, etc.

References

- AWS Certified Advanced Networking Specialty
- Advanced Networking Exam Details
- AWS Global Transit Network
- Youtube: Re:Invent: Connecting many VPCs
- AWS: Multiple Region, Multiple VPC Connectivity
- AWS: VPN Connection User Guide
- AWS Whitepapers