Well-Architected Direct Connect

Steven Moran
TECHNICAL INSTRUCTOR

Resiliency – Multiple Connections

- Implement multiple DX connections to increase resiliency.
  - Multiple connections at single DX location.
  - Multiple DX Locations.
- Implement LAGs.
- Direct Connect Resiliency Toolkit

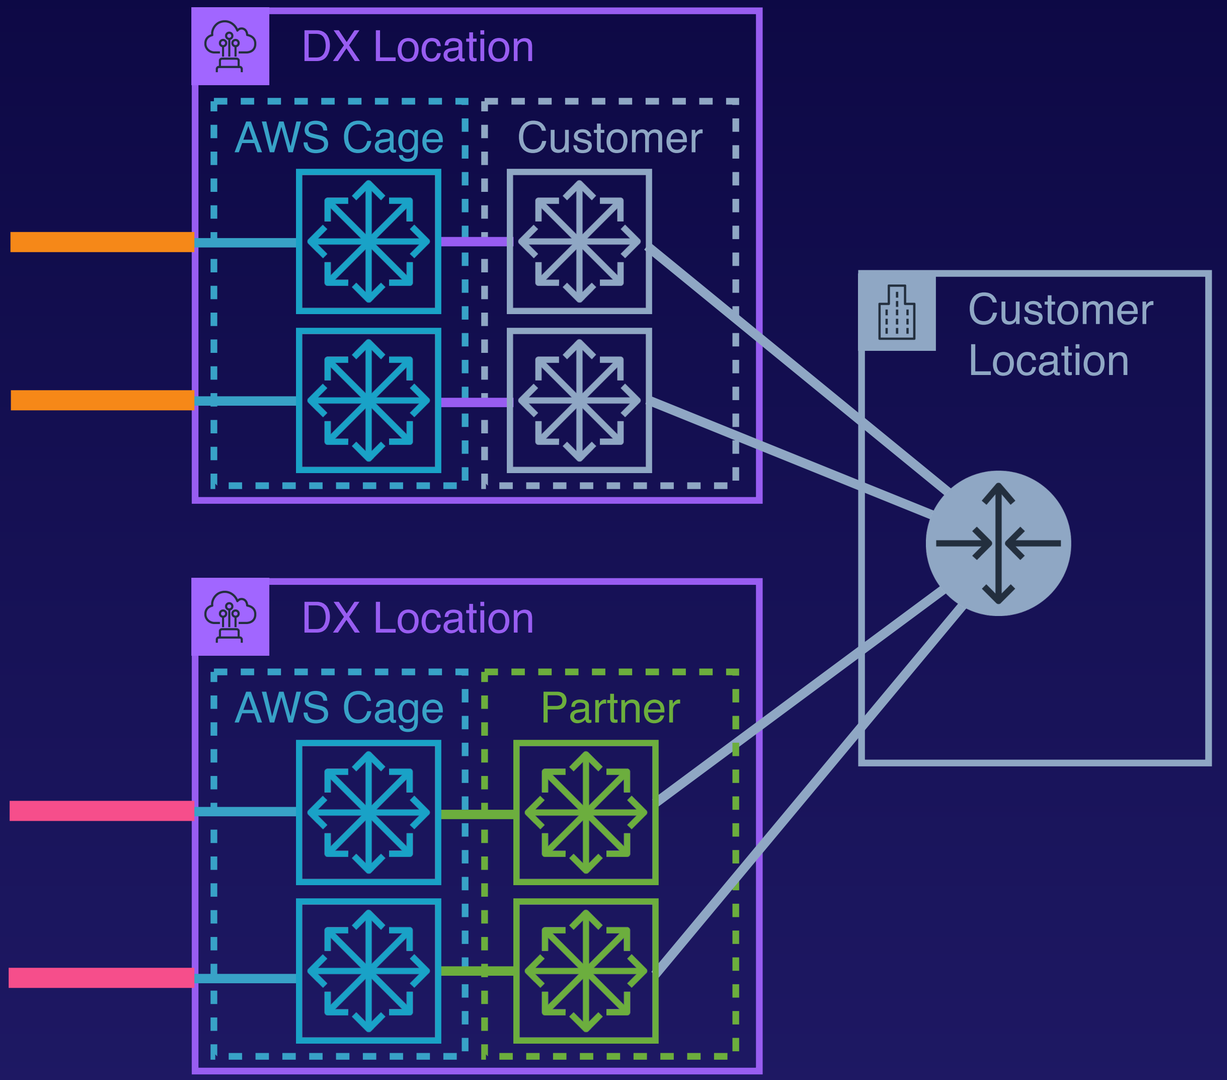

Resiliency – Multiple Connections

Resiliency – Multiple Connections

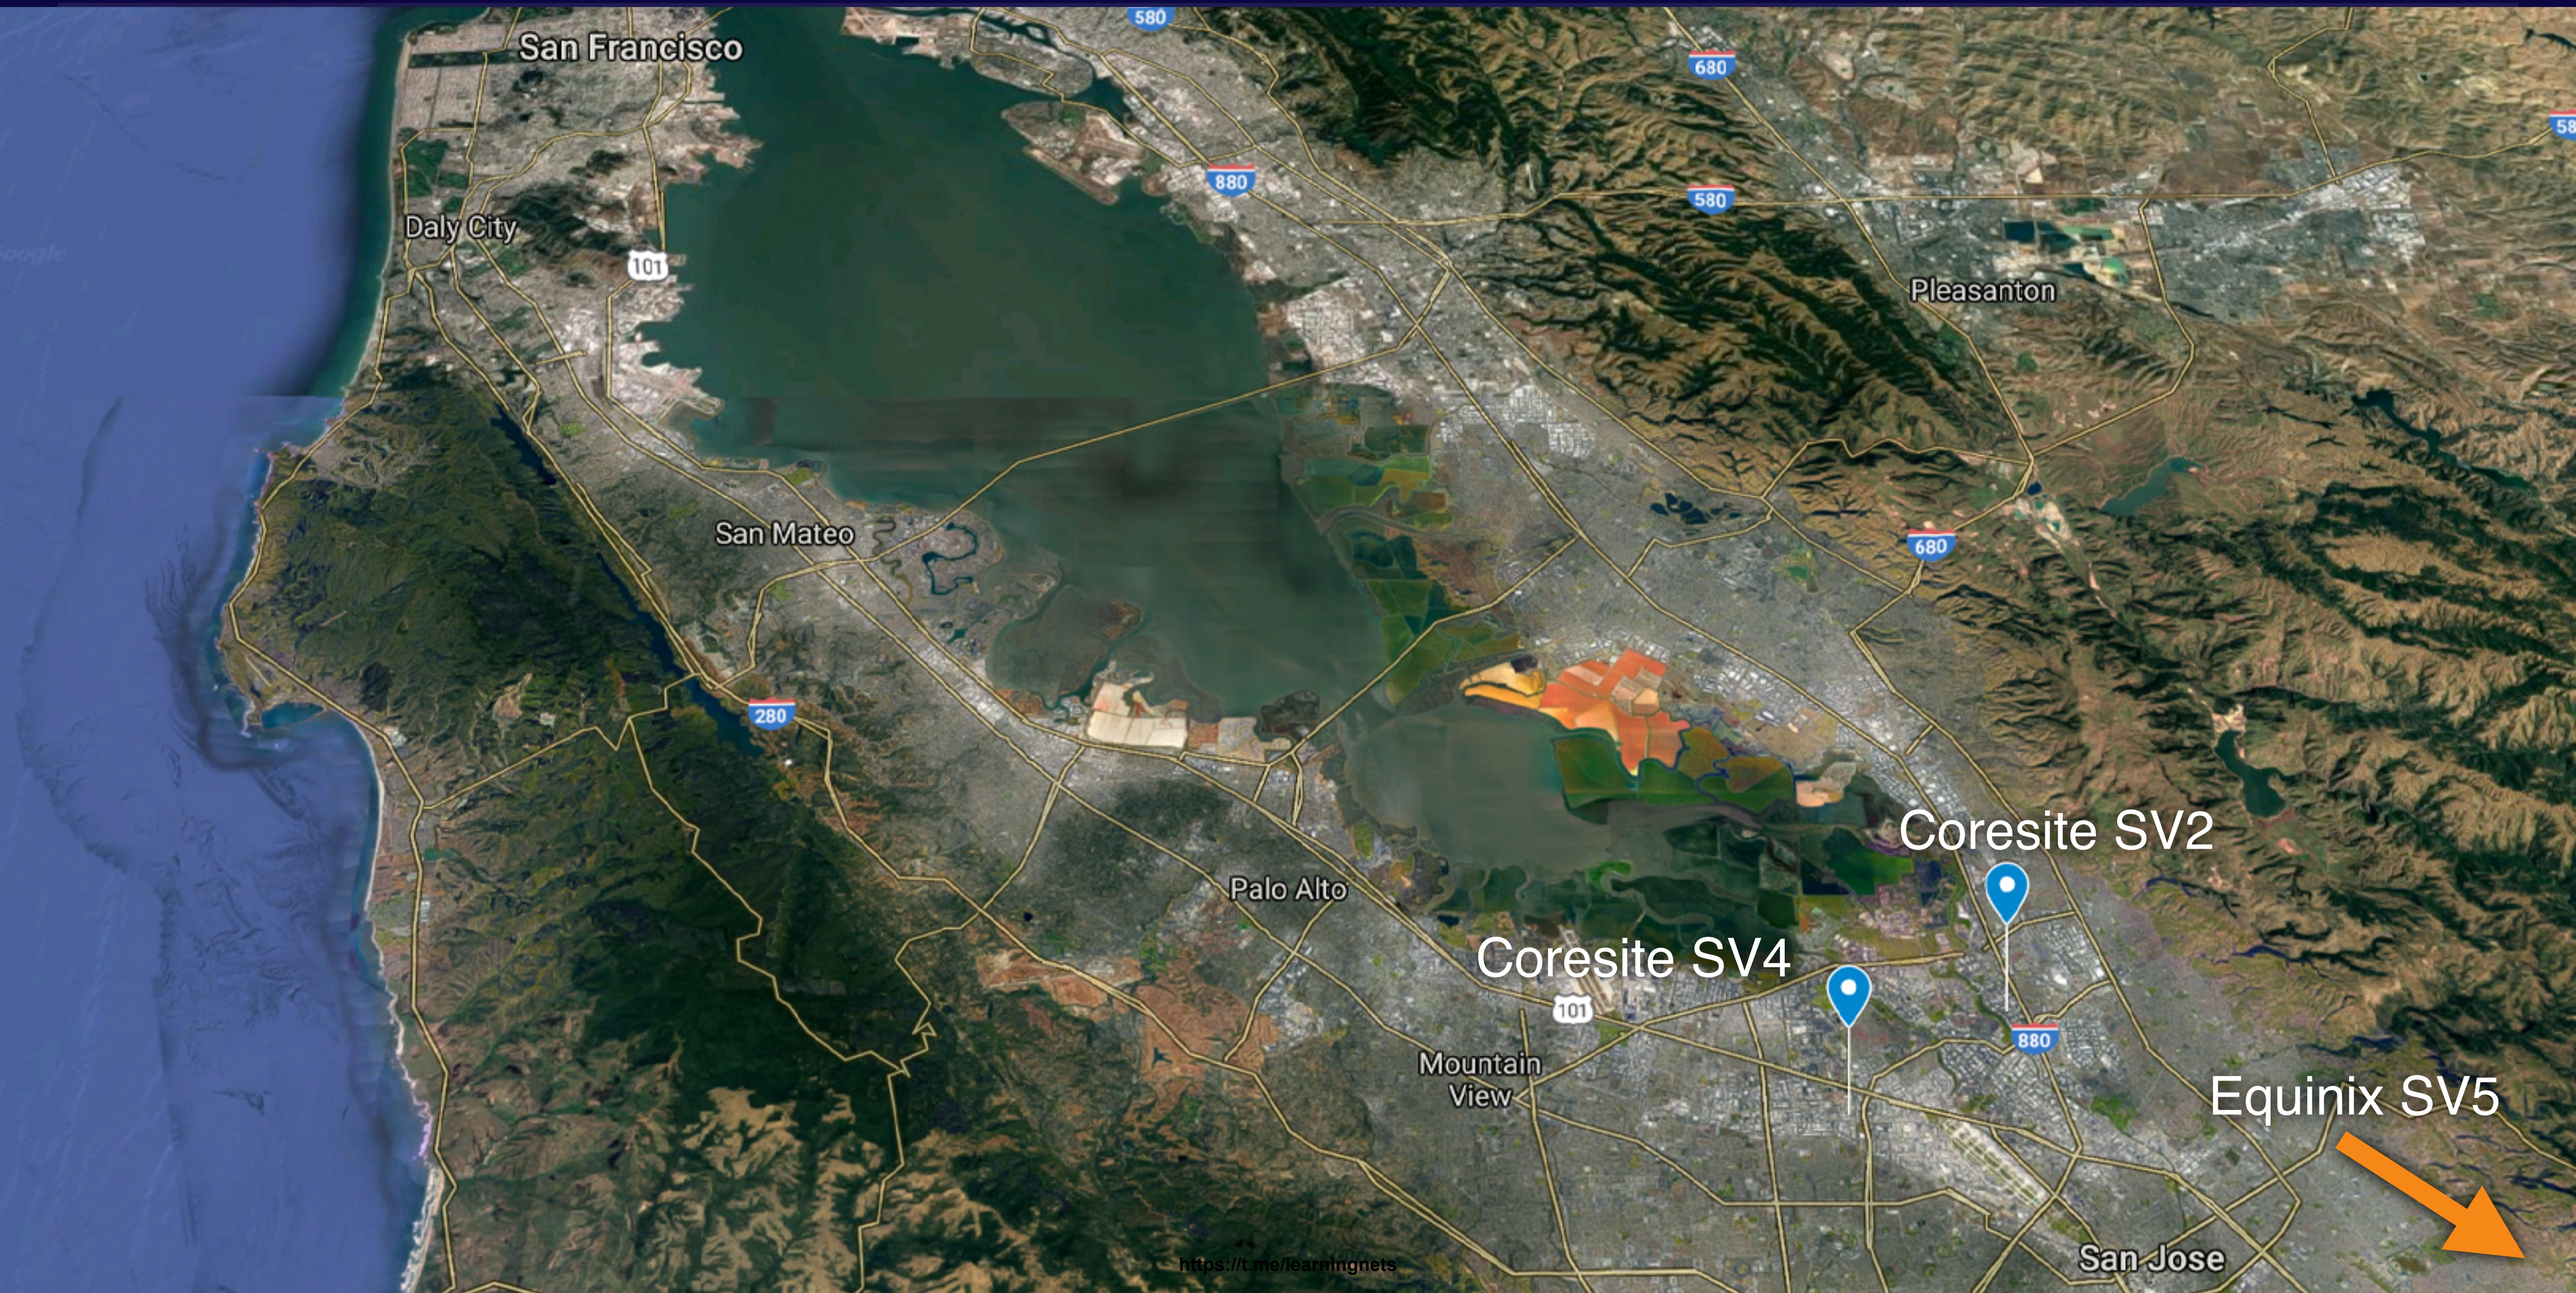

Resiliency – Multiple Connections

Resiliency – Multiple Connections

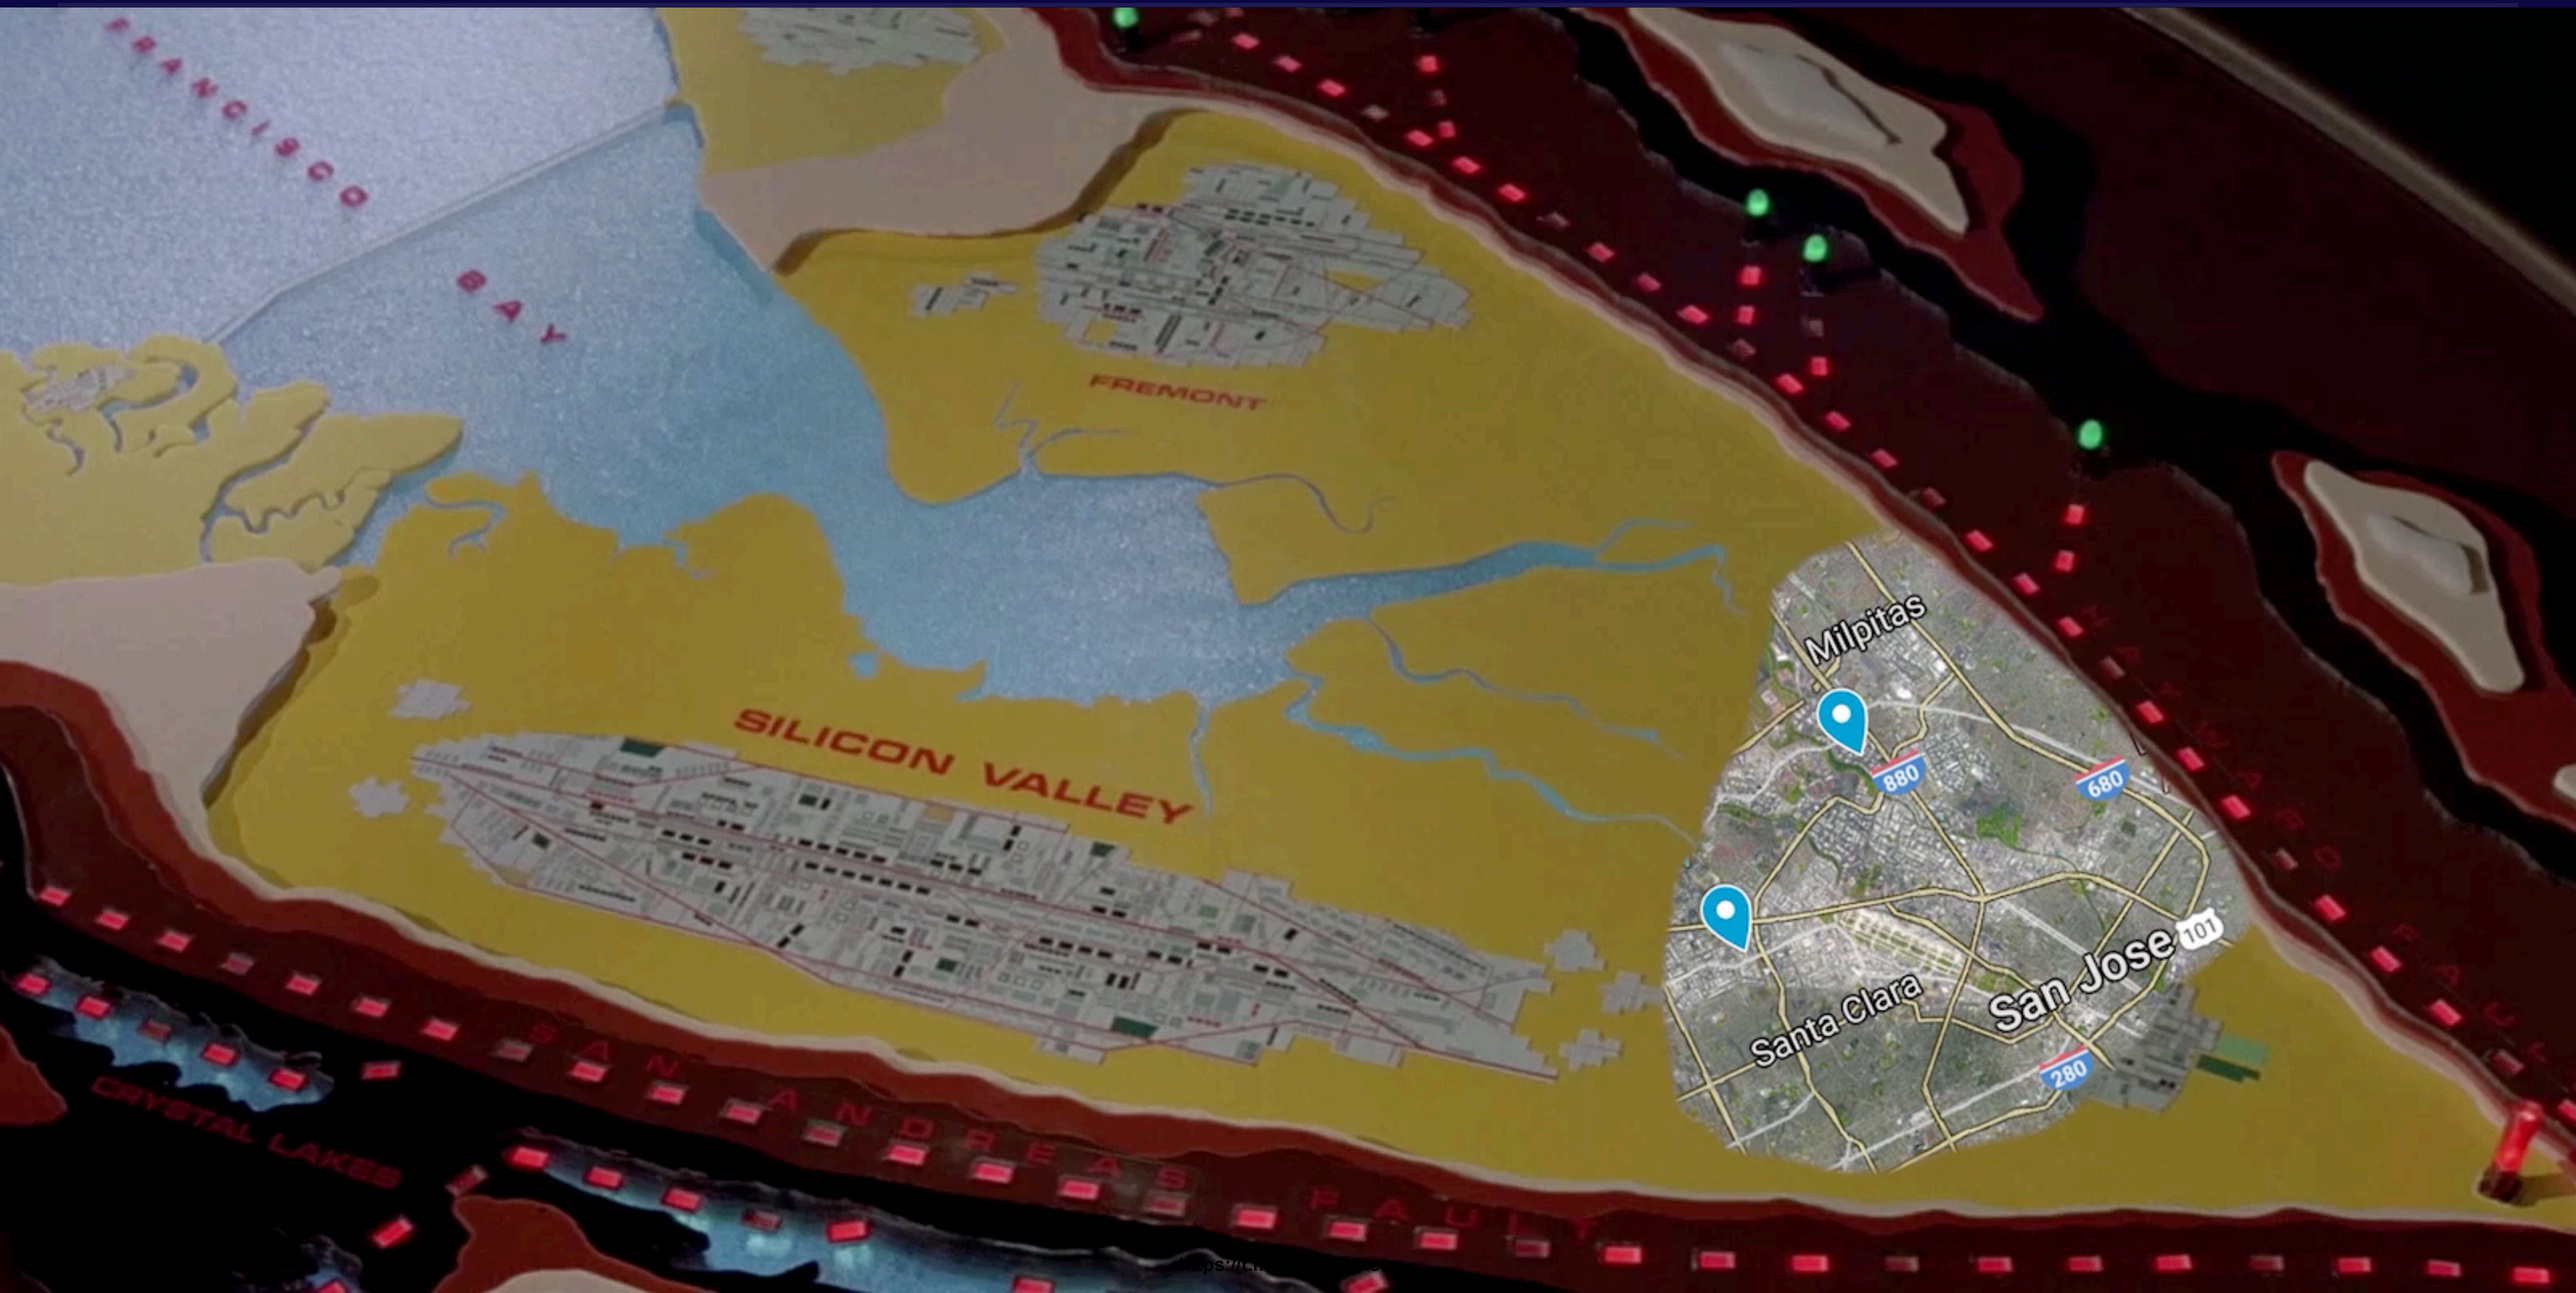

Resiliency – Bidirectional Forwarding Detection

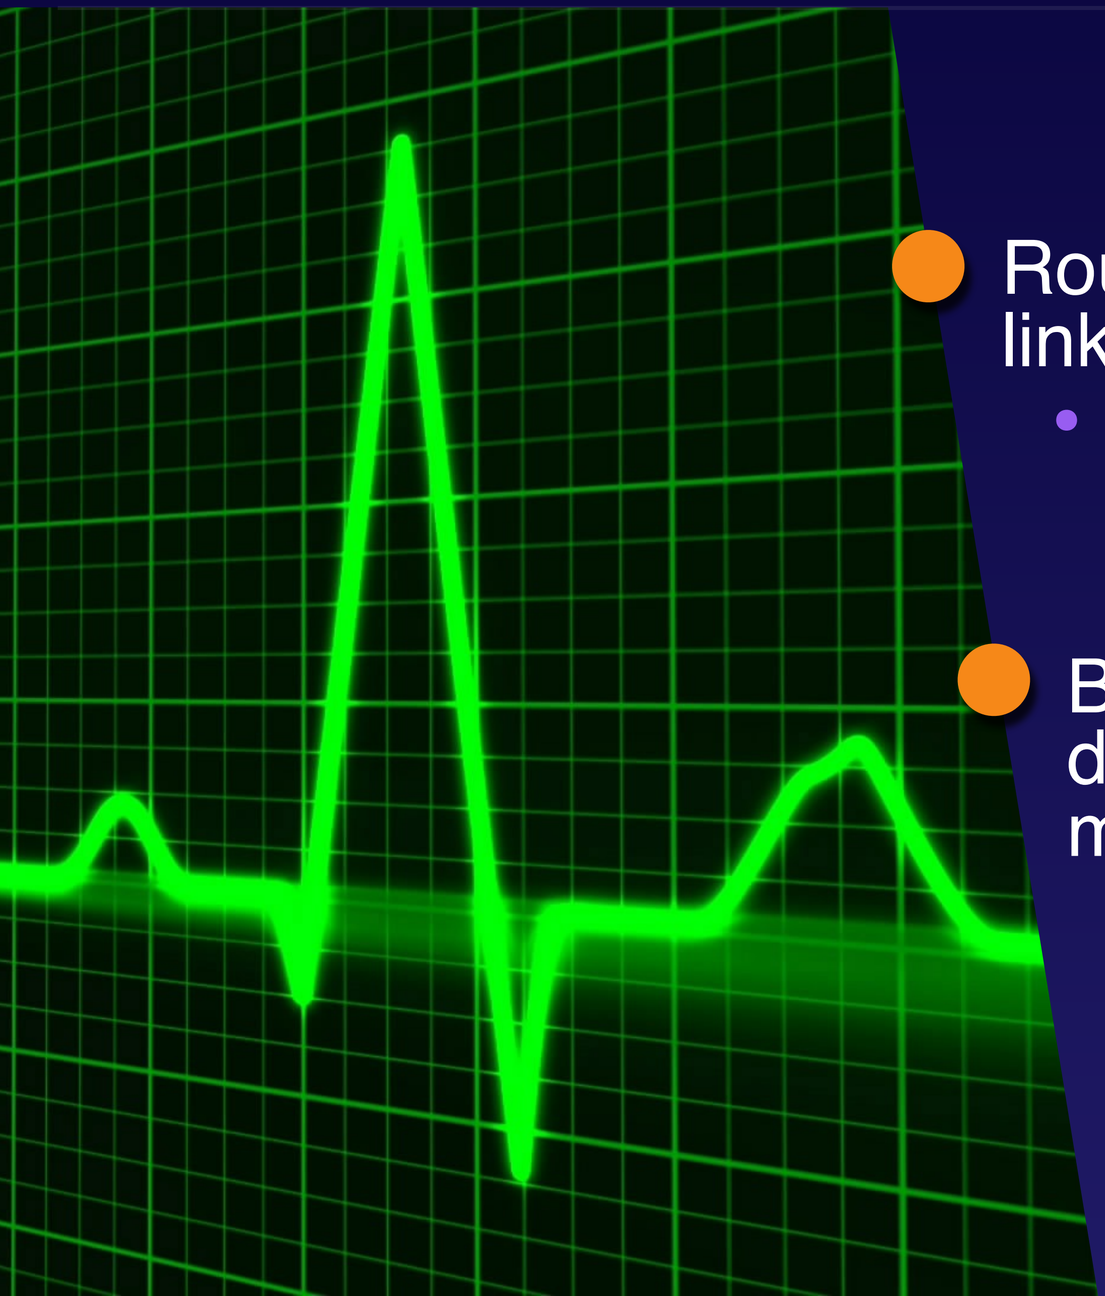

- Routing protocols do not quickly recognize link failures.
  - BGP default peer timeout is 270 seconds.
- Bidirectional Forwarding Detection can detect link failures and respond within milliseconds.
  - Informs associated routing protocol to recalculate best paths.

Resiliency – Bidirectional Forwarding Detection

- Automatically enabled on VIFs.
- Must be configured at customer devices.
  - Tx/Rx intervals in milliseconds.
  - Interval multiplier.

Resiliency – Failover to VPN

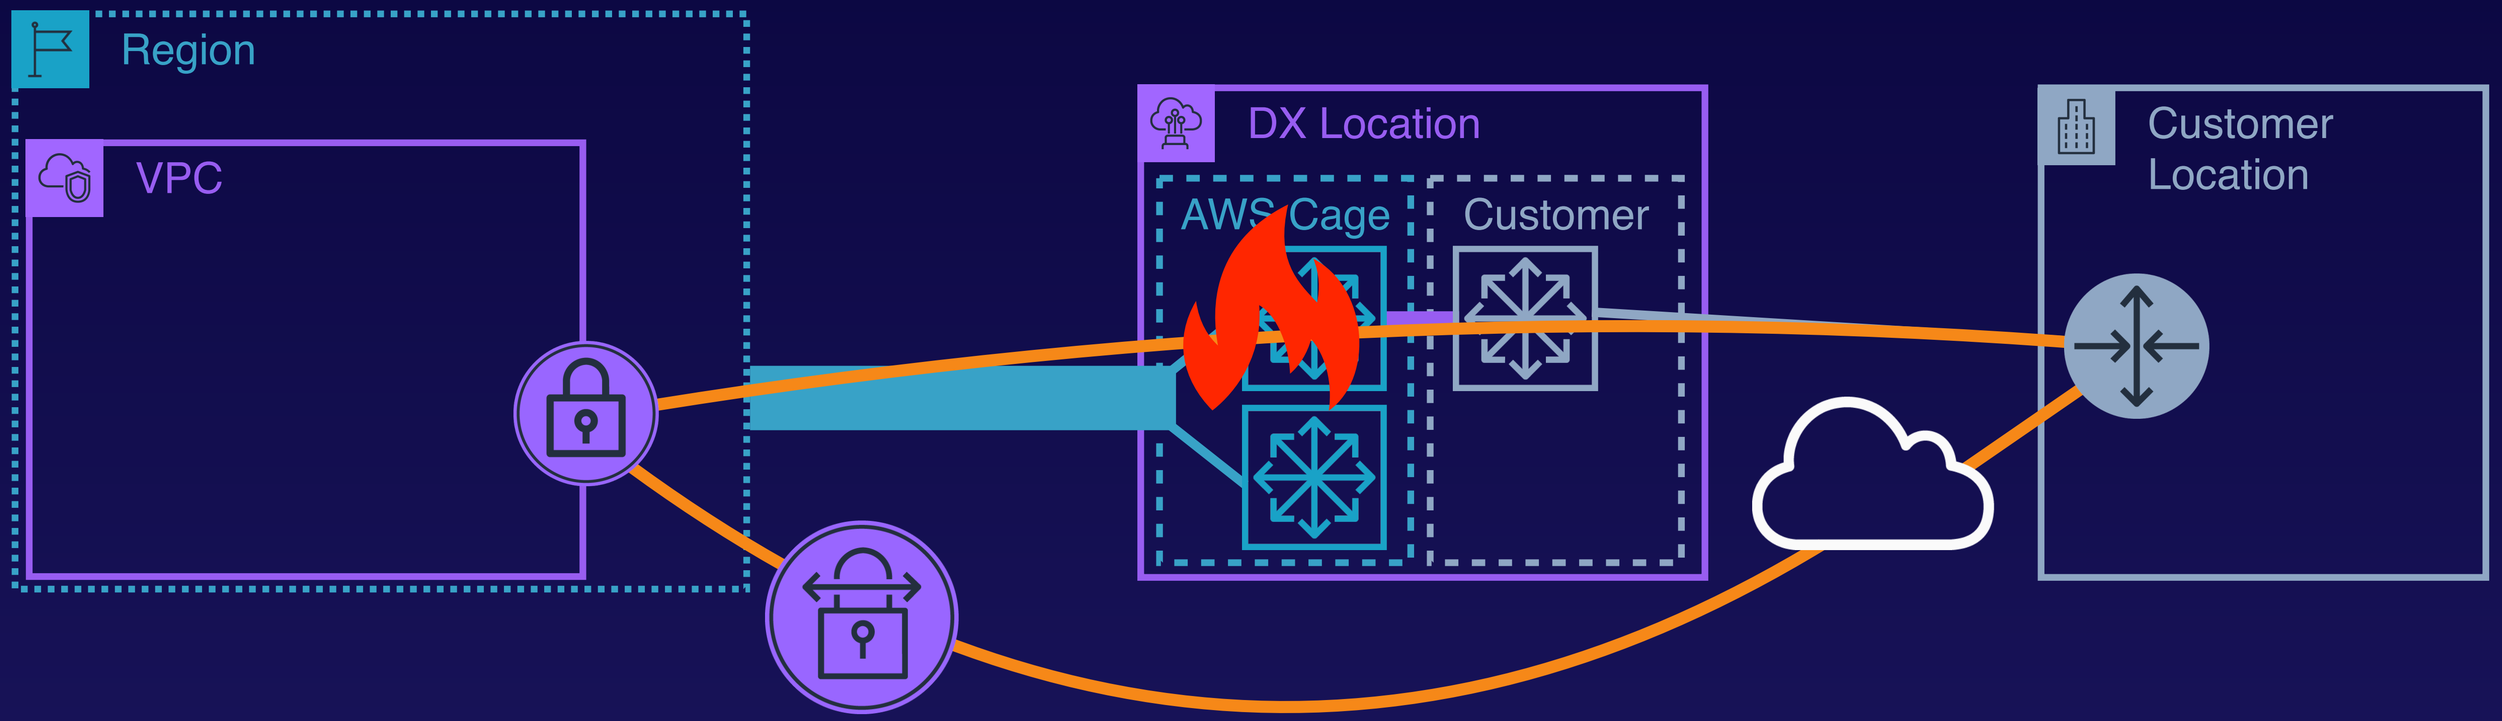

- BGP prefix preference configuration can be used to initiate VPN connections if DX paths are down.
- AWS VPN creation can be automated.
- Create VPN server AMIs.

Resiliency & Performance – Local BGP Communities

Control prefix-preference for AWS traffic returning to your network.

Used with multiple DX connections to configure load balancing or failover.

Supported over Private and Transit VIFs.

Resiliency & Performance – Local BGP Communities

A CLOUD GURU

Community values:

- Low preference – 7224:7100
- Medium preference – 7224:7200
- High preference – 7224:7300

Traffic will be sent to higher-preference prefixes over lower-preference prefixes.

Traffic is load balanced across multiple, equal-preference prefixes.

Performance – LAGs

- LAGs can be implemented to aggregate throughput from multiple connections.
- 4 x 10 Gbps physical connections = 1 x 40 Gbps logical connection.

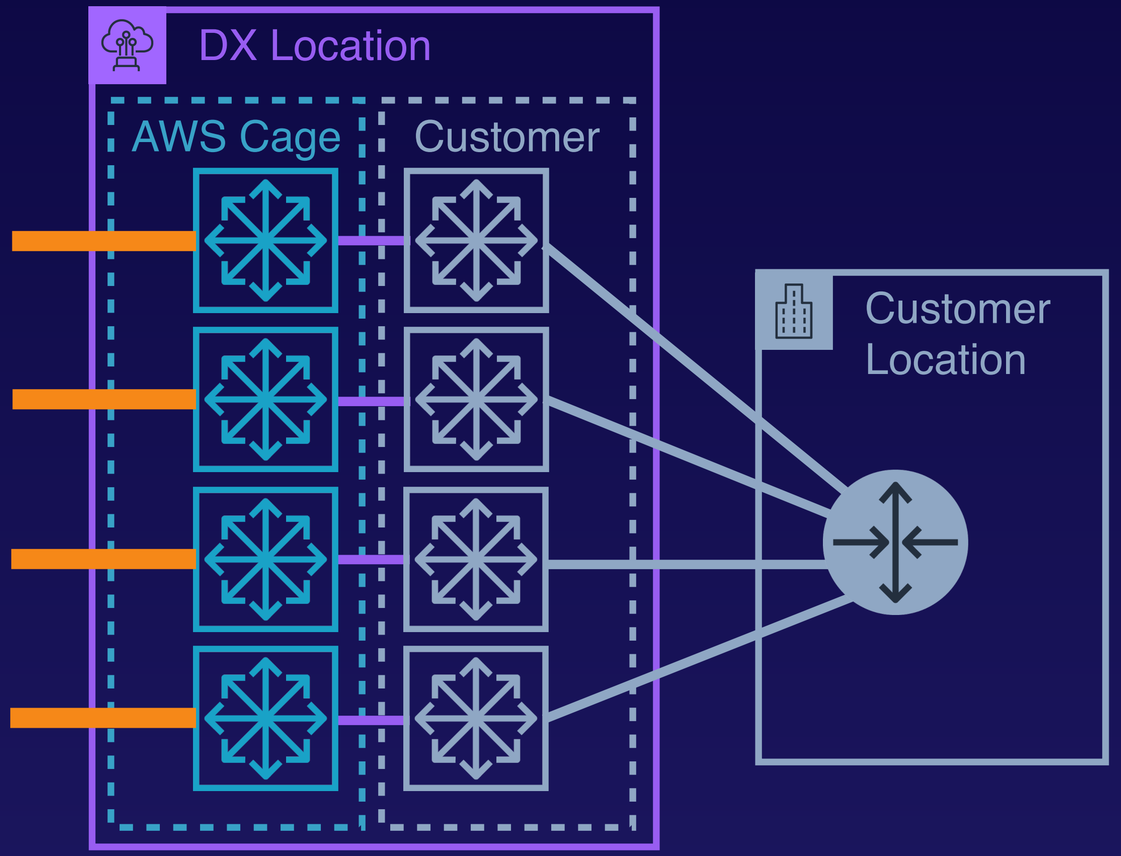

Performance – Jumbo Frames

- Enabling jumbo frames for private VIFs sets MTU from 1,500 to 9,001 bytes.
  - More data sent per Ethernet packet.
  - Total data transmitted with fewer packets.
- Enabling/disabling jumbo frames on a VIF disrupts traffic for all VIFs on the connection for up to 30 seconds.

Performance – Jumbo Frames

- Ensure that support for jumbo frames is enabled from end to end.
- Using nested VLANs might cause occasional data corruption if jumbo frames are not enabled.
- Not not supported via VPG static routes.

Security – VPN over DX

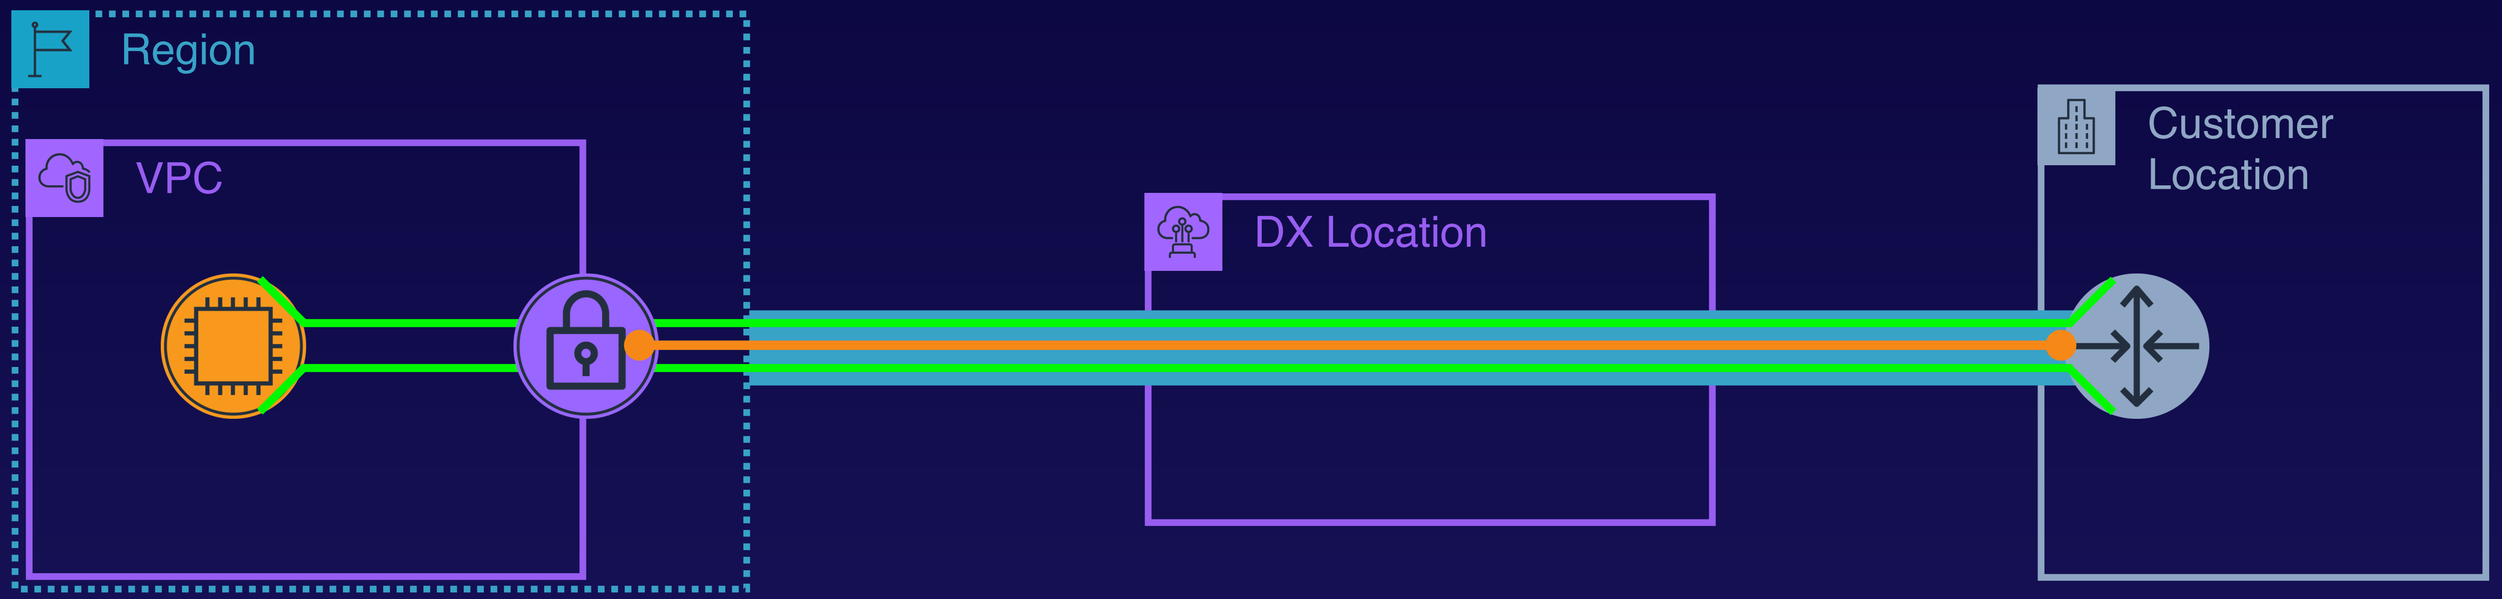

- AWS does not encrypt DX traffic.
- VPN tunnels may be established from on-prem to EC2-based VPN systems.

Cost Optimization - DX DTO vs Internet DTO

Direct Connect has two billing elements:

- Port Hours
- Data Transfer Out (DTO)

DTO rates vary depending on the source AWS Region and the DX location the traffic is sent to.

DX DTO rate is not reduced by increased usage.

Data transferred in is always free.

Cost Optimization - DX DTO vs Internet DTO Scenario

For “x” amount of DTO, when does DX become cheaper than internet DTO?

- Scenario conditions:
  - Internet DTO rate: $0.05/GB
  - DX DTO rate: $0.02/GB
  - 1G Port: $0.30/hour
  - 10G Port: $2.25/hour

(.05x) - ((.02x) + (port rate * time))

Negative amounts mean that DX is more expensive.

Internet DTO vs DX DTO Price Delta – 1G Port

A CLOUD GURU

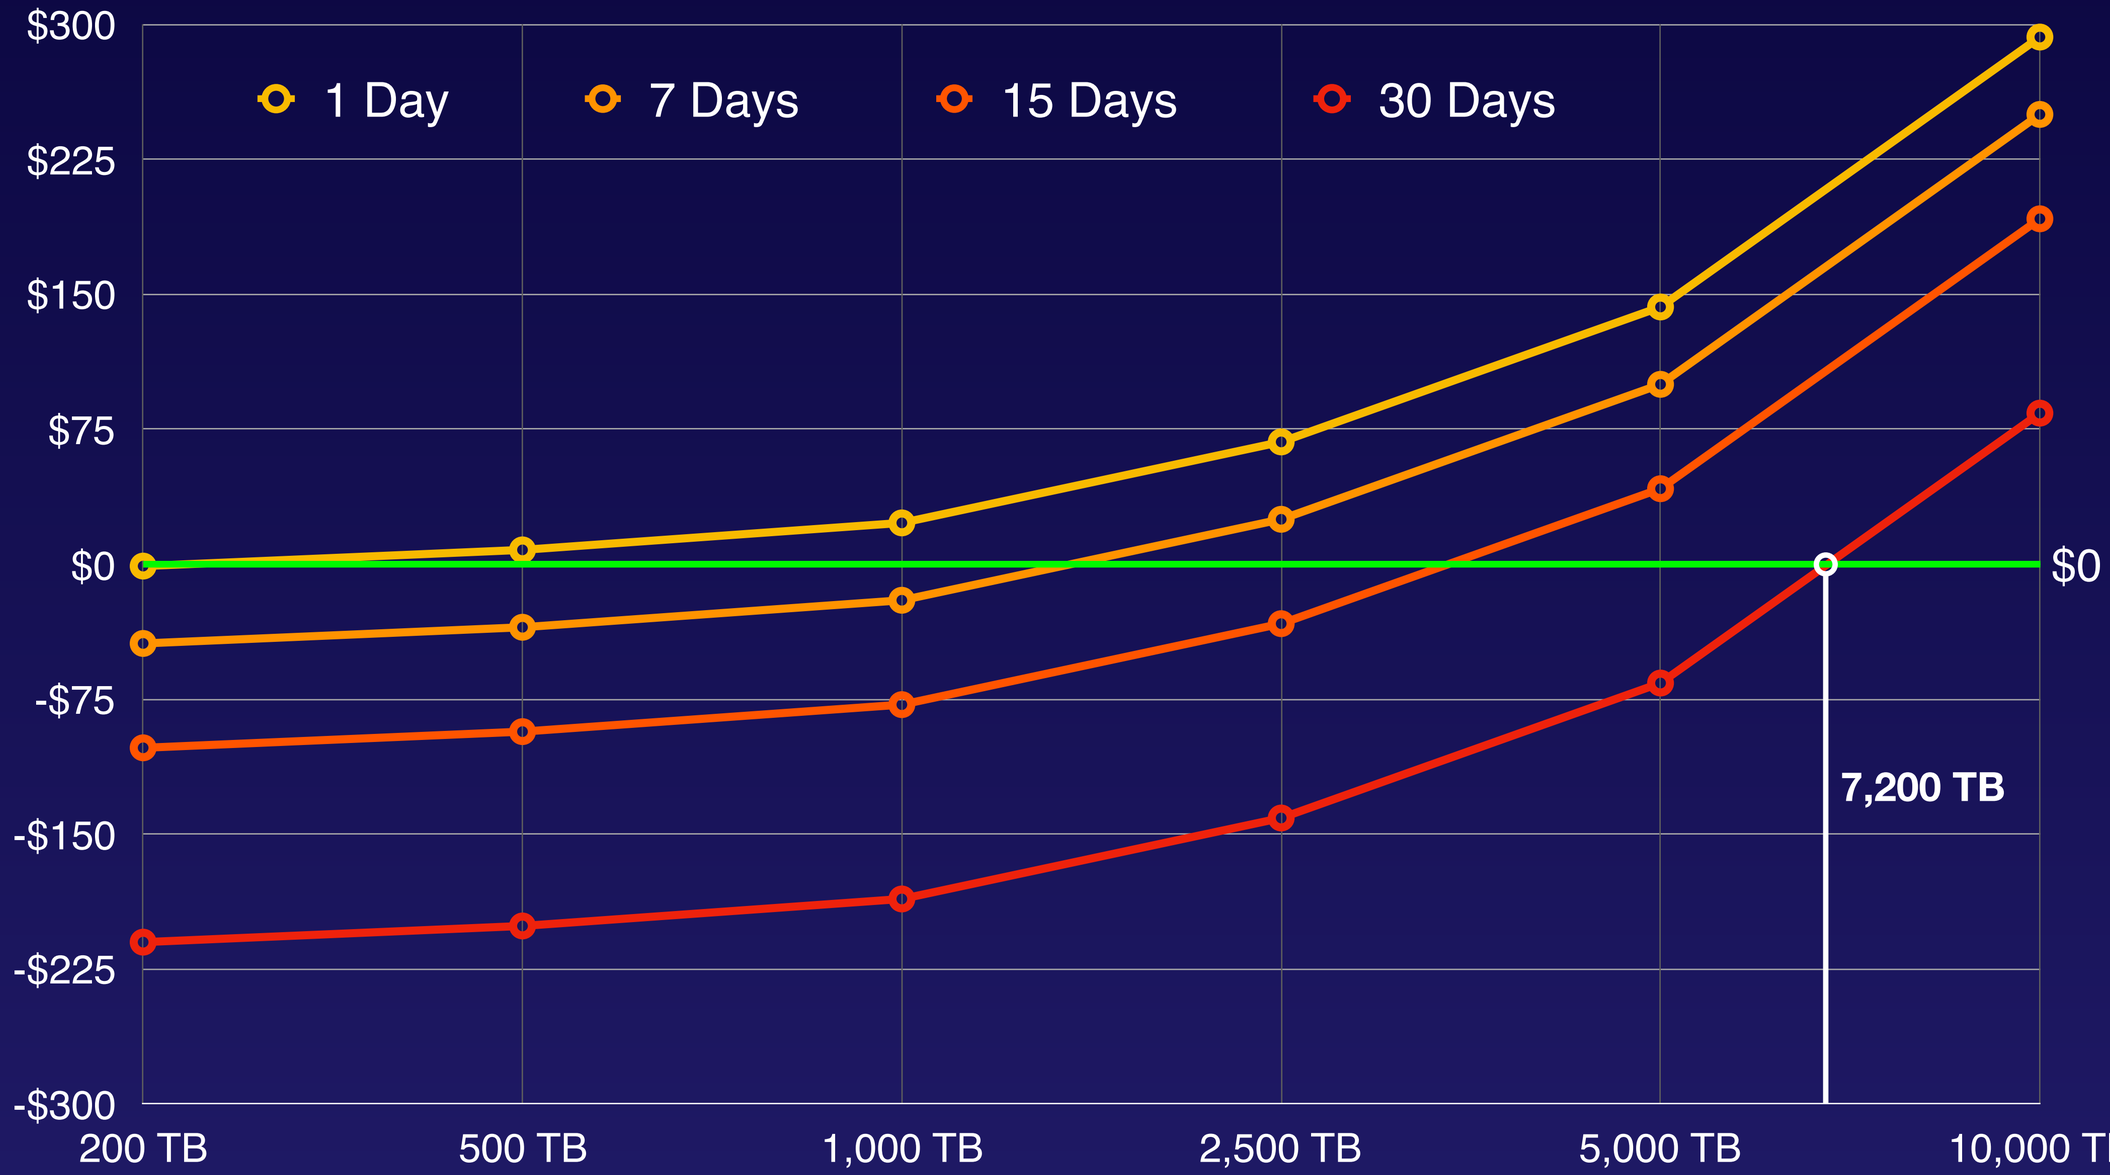

Internet DTO vs DX DTO Price Delta - 10G Port

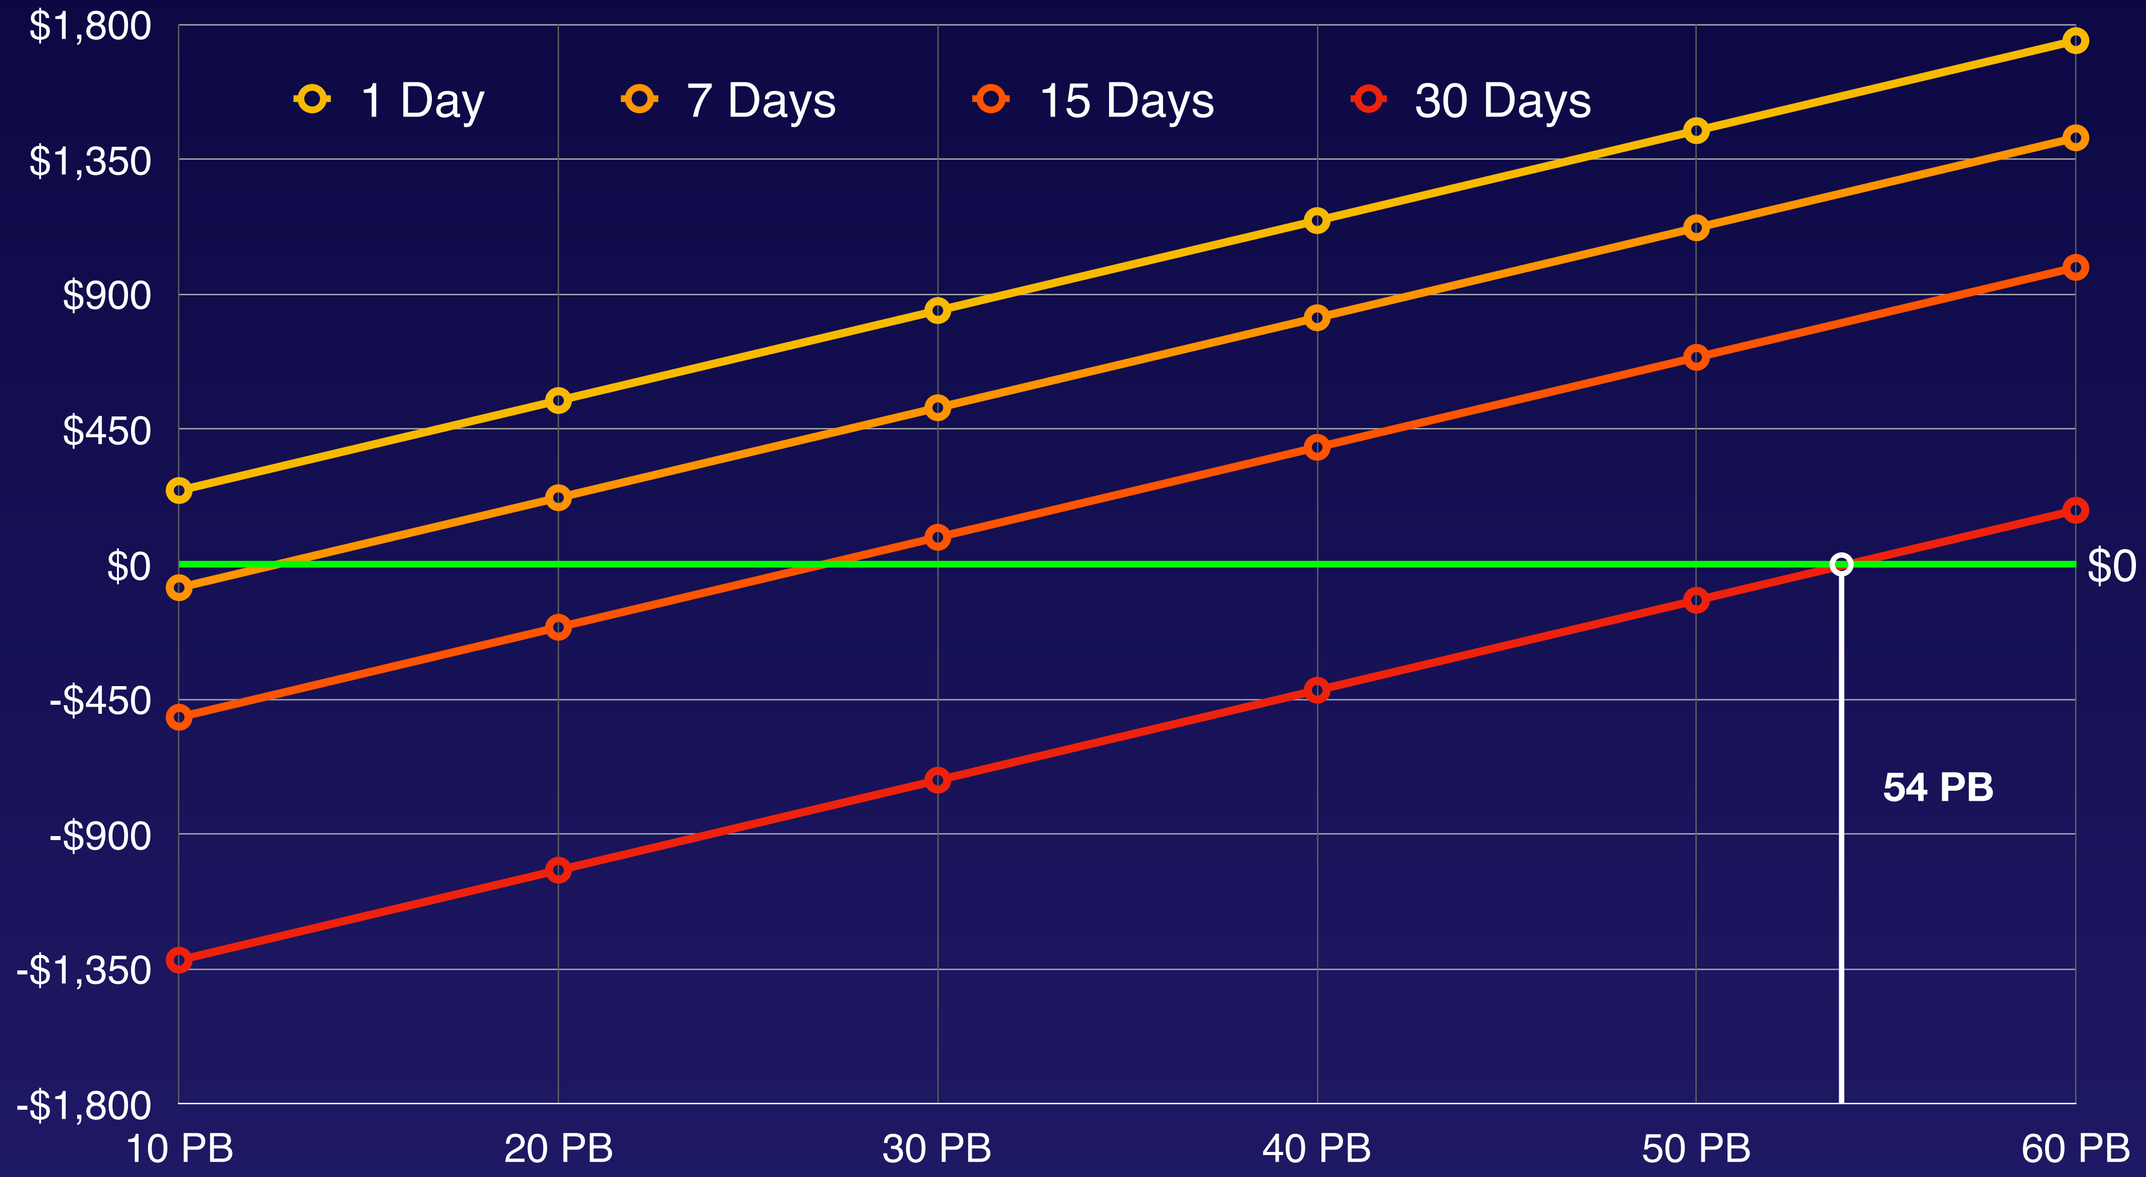

Operational Excellence – Monitoring

- CloudWatch metrics for DX connections:
  - ConnectionState
  - ConnectionBpsEgress/Ingress
  - ConnectionPpsEgress/Ingress
  - ConnectionCRCErrorCount
  - ConnectionLightLevelTx/Rx
- VIFs do not have CW metrics.
  - VIF status must be checked manually.

Operational Excellence – Troubleshooting

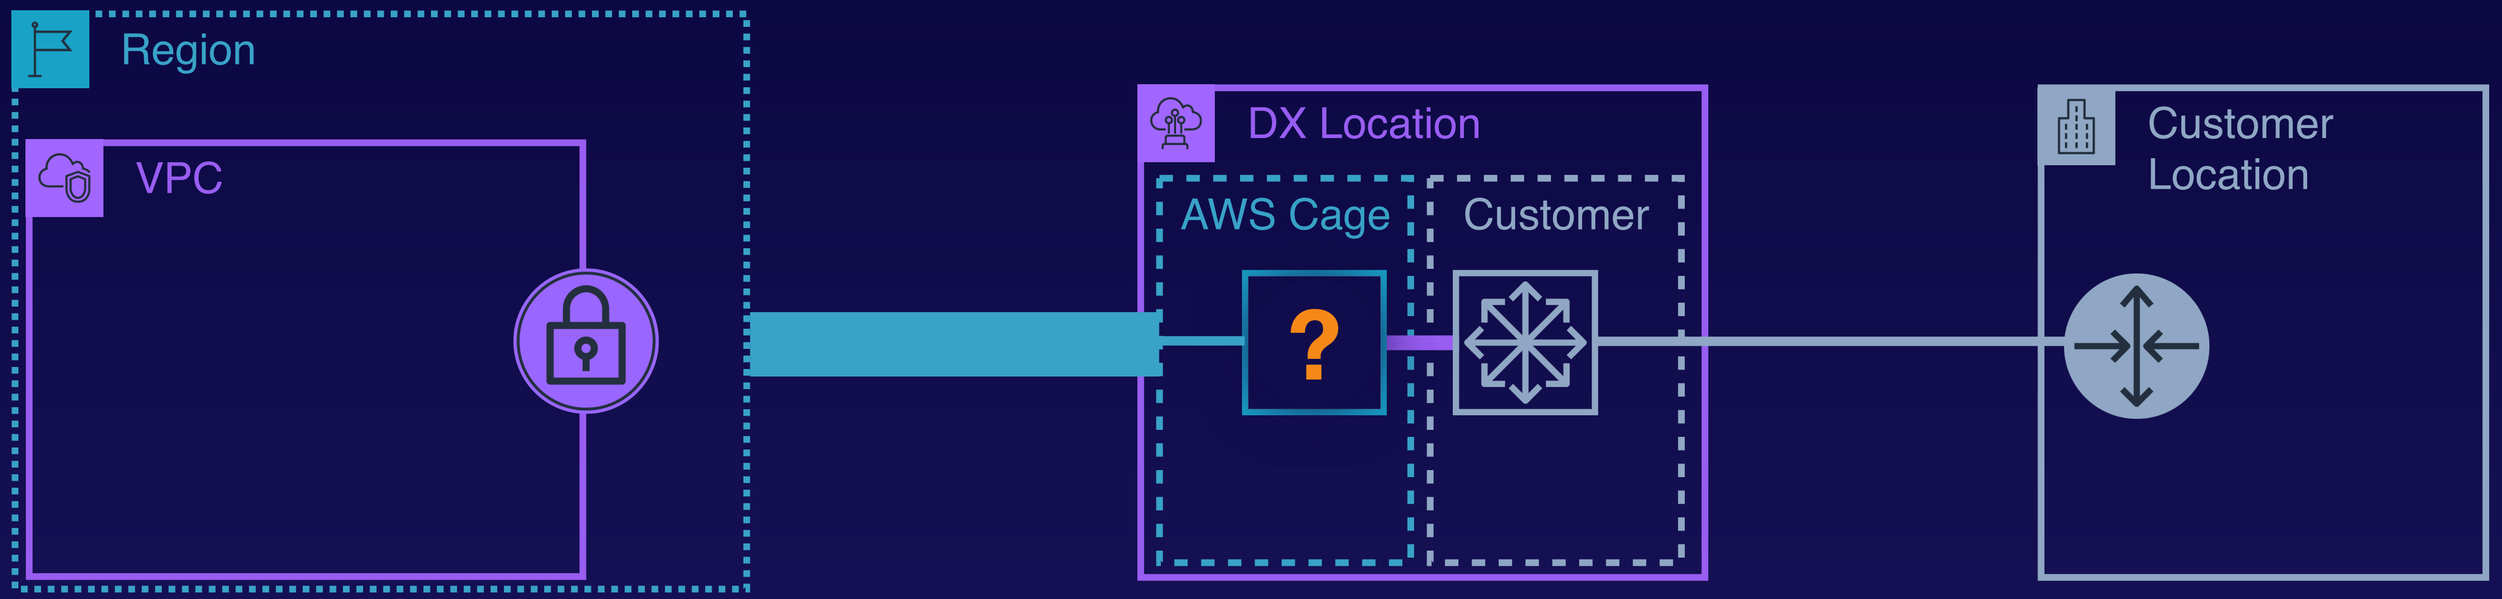

- No connectivity to AWS device specified in LOA-CFA?

Operational Excellence – Troubleshooting

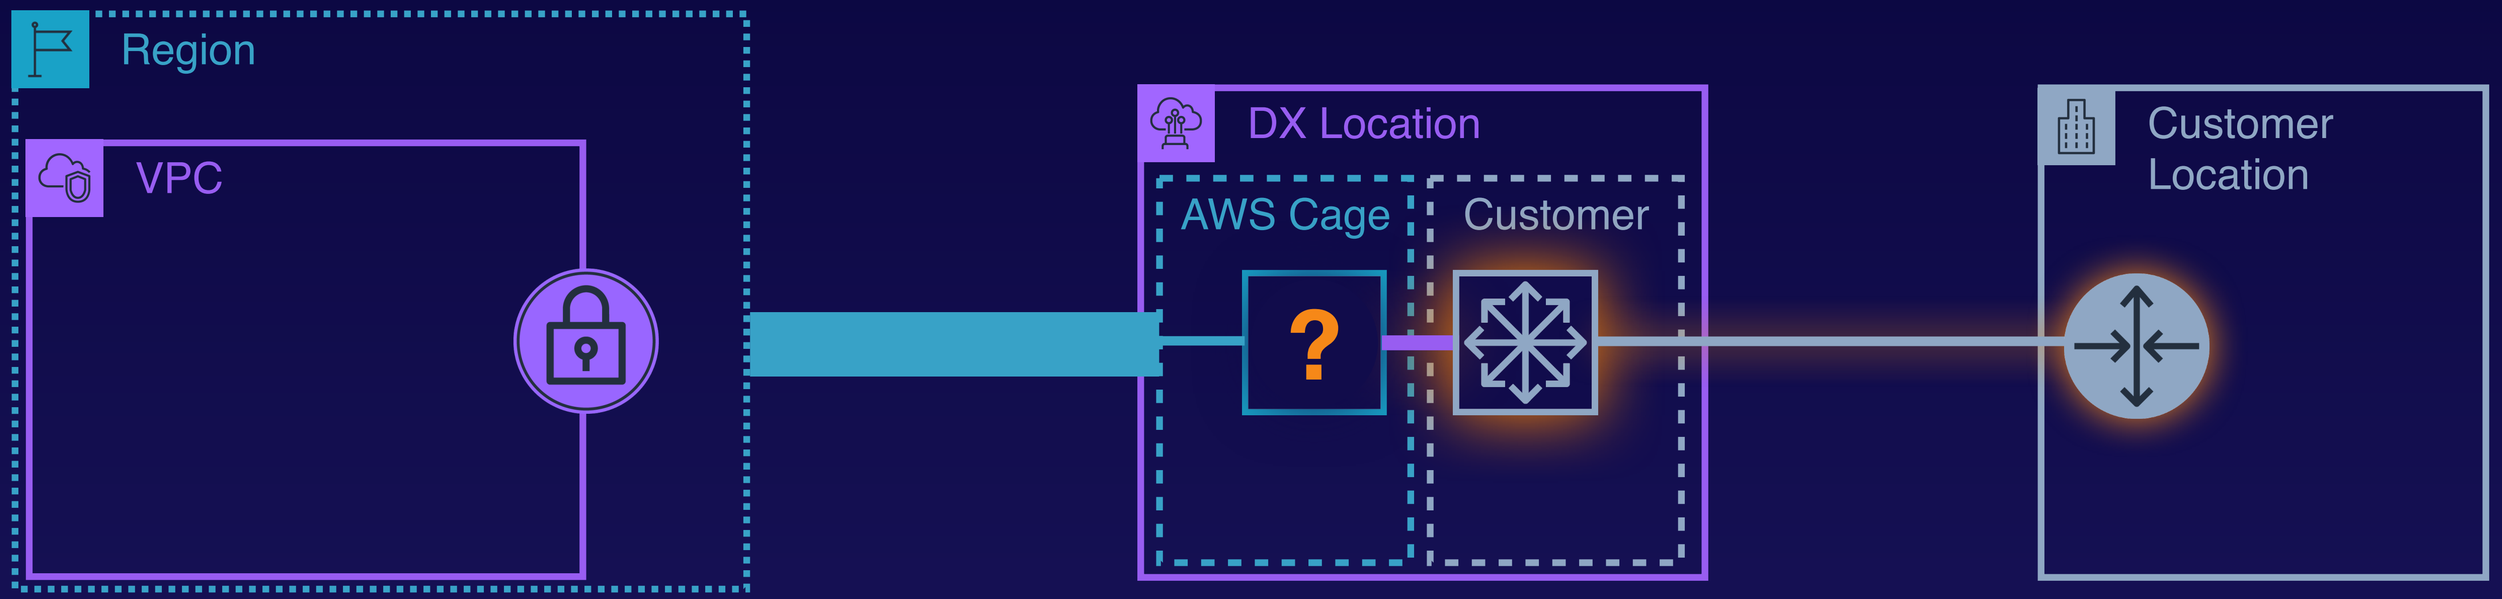

- No connectivity to AWS device specified in LOA-CFA?
  - Physical layer problem at on-prem or DX location.
  - Check cable and port configuration.
  - Ensure DX hardware requirements are followed.

Operational Excellence – Troubleshooting

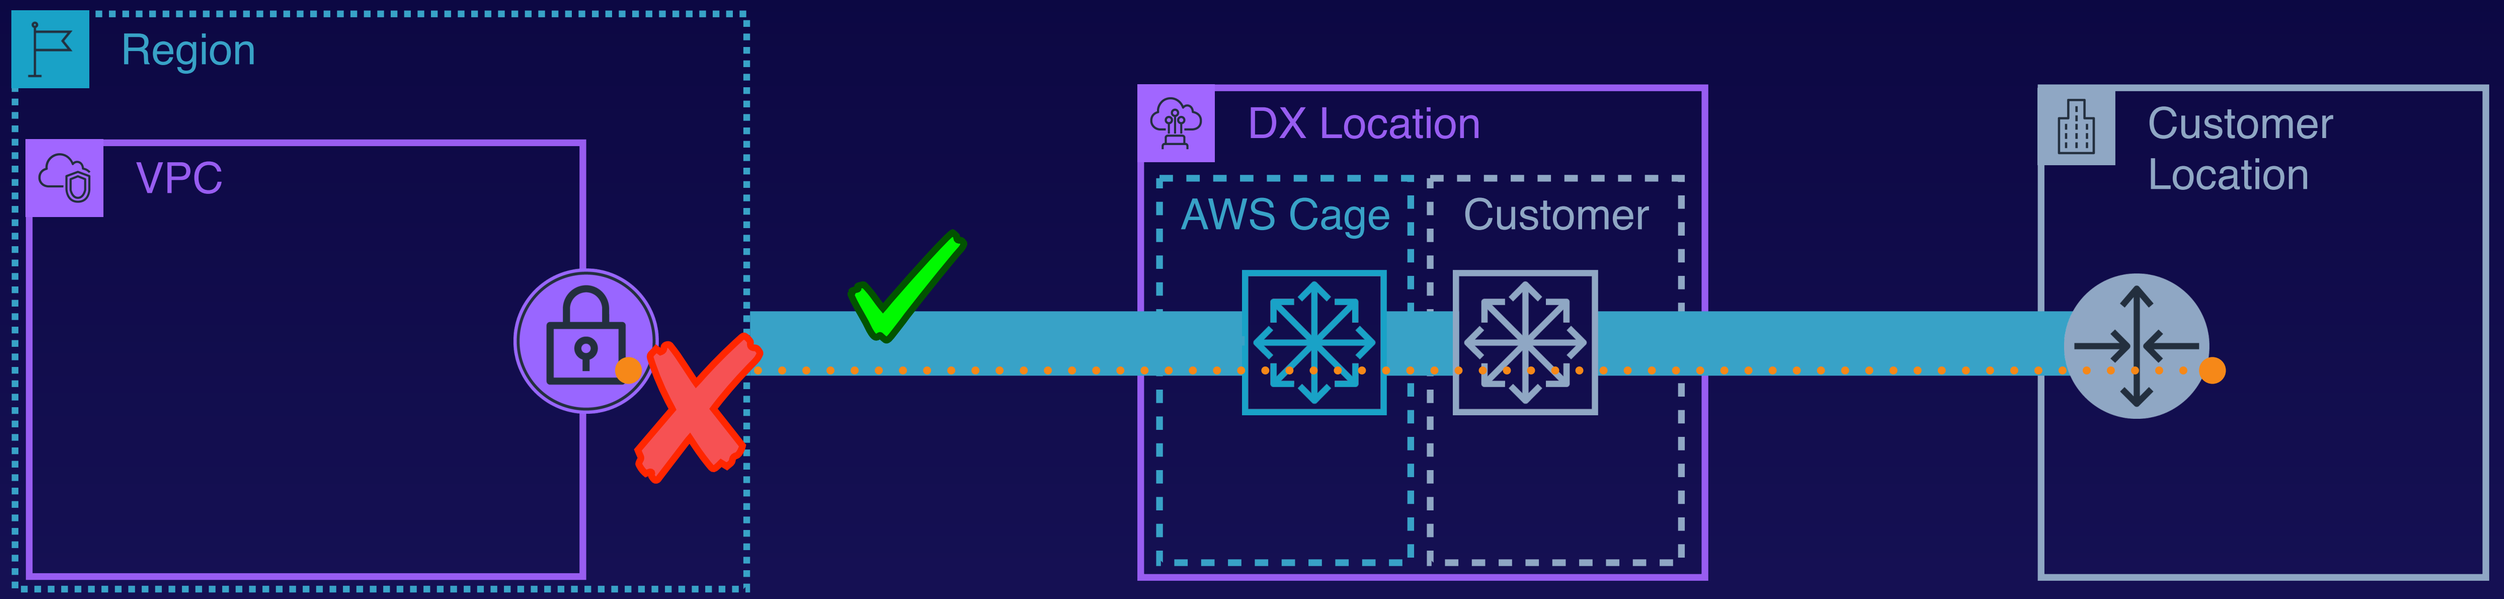

- DX connection is up but VIFs stay down?
  - Data link layer problem.

Operational Excellence – Troubleshooting

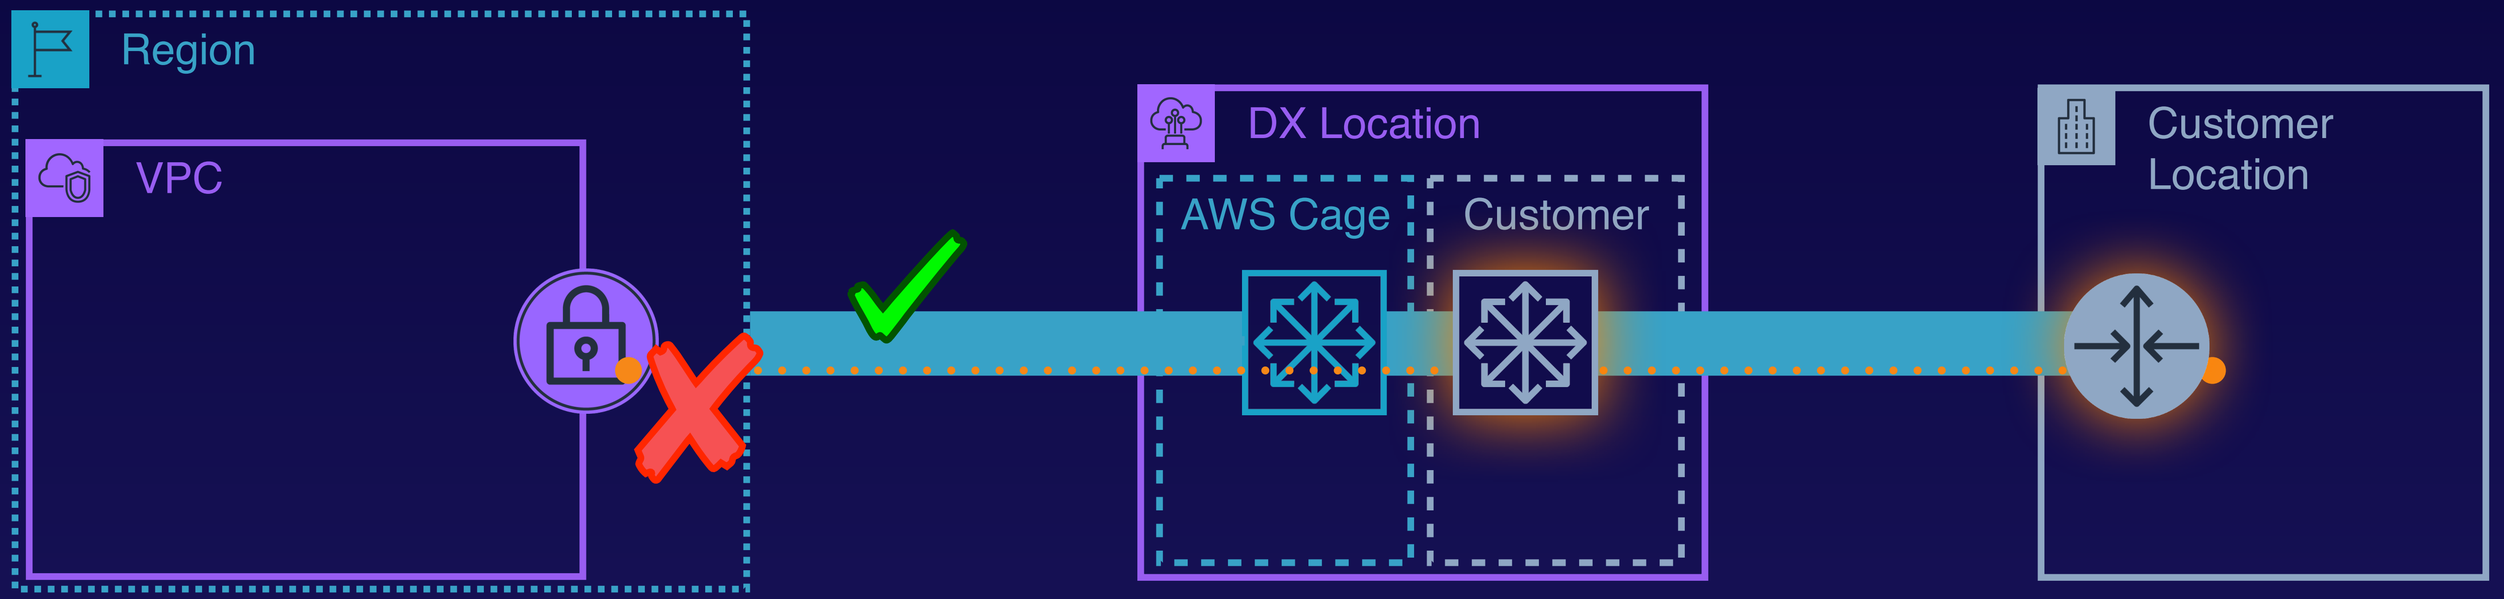

- DX connection is up but VIFs stay down?
  - Data link layer problem.
  - Verify 802.1Q support.
  - Ensure correct VLAN and VIF configuration.

Operational Excellence – Troubleshooting

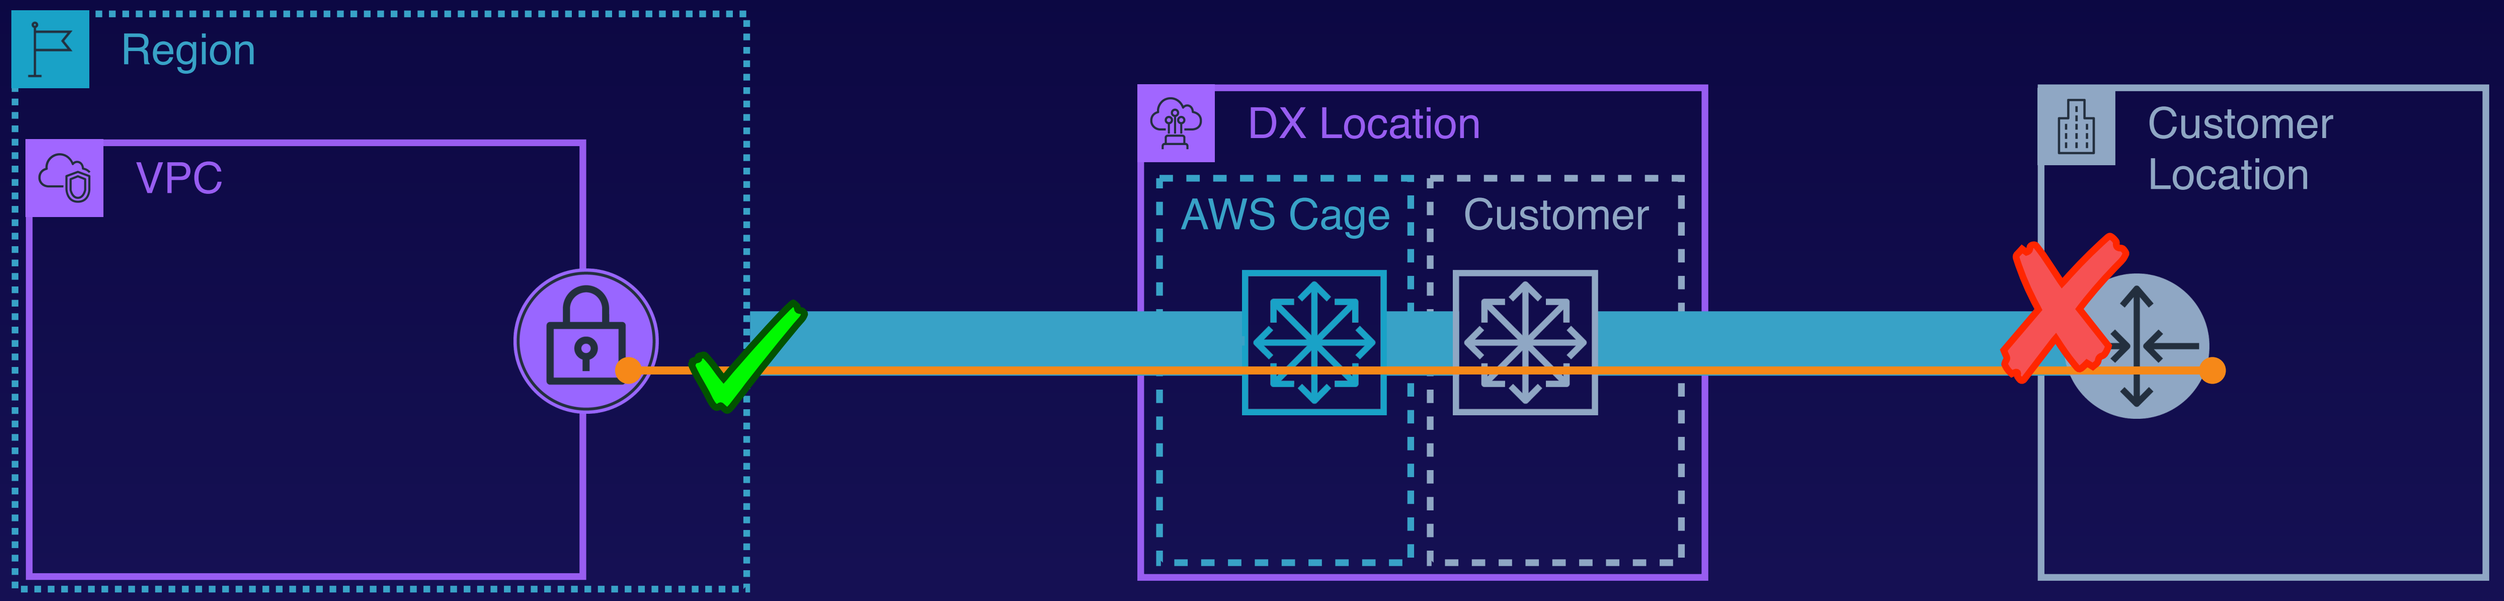

- Connection up but can't establish BGP session?
- Misconfigured:
  - ASNs
  - Peer IP address
  - MD5 authentication key
  - Exceeded max number of advertised prefixes
  - BGP port (TCP 179) is blocked

Operational Excellence – Troubleshooting

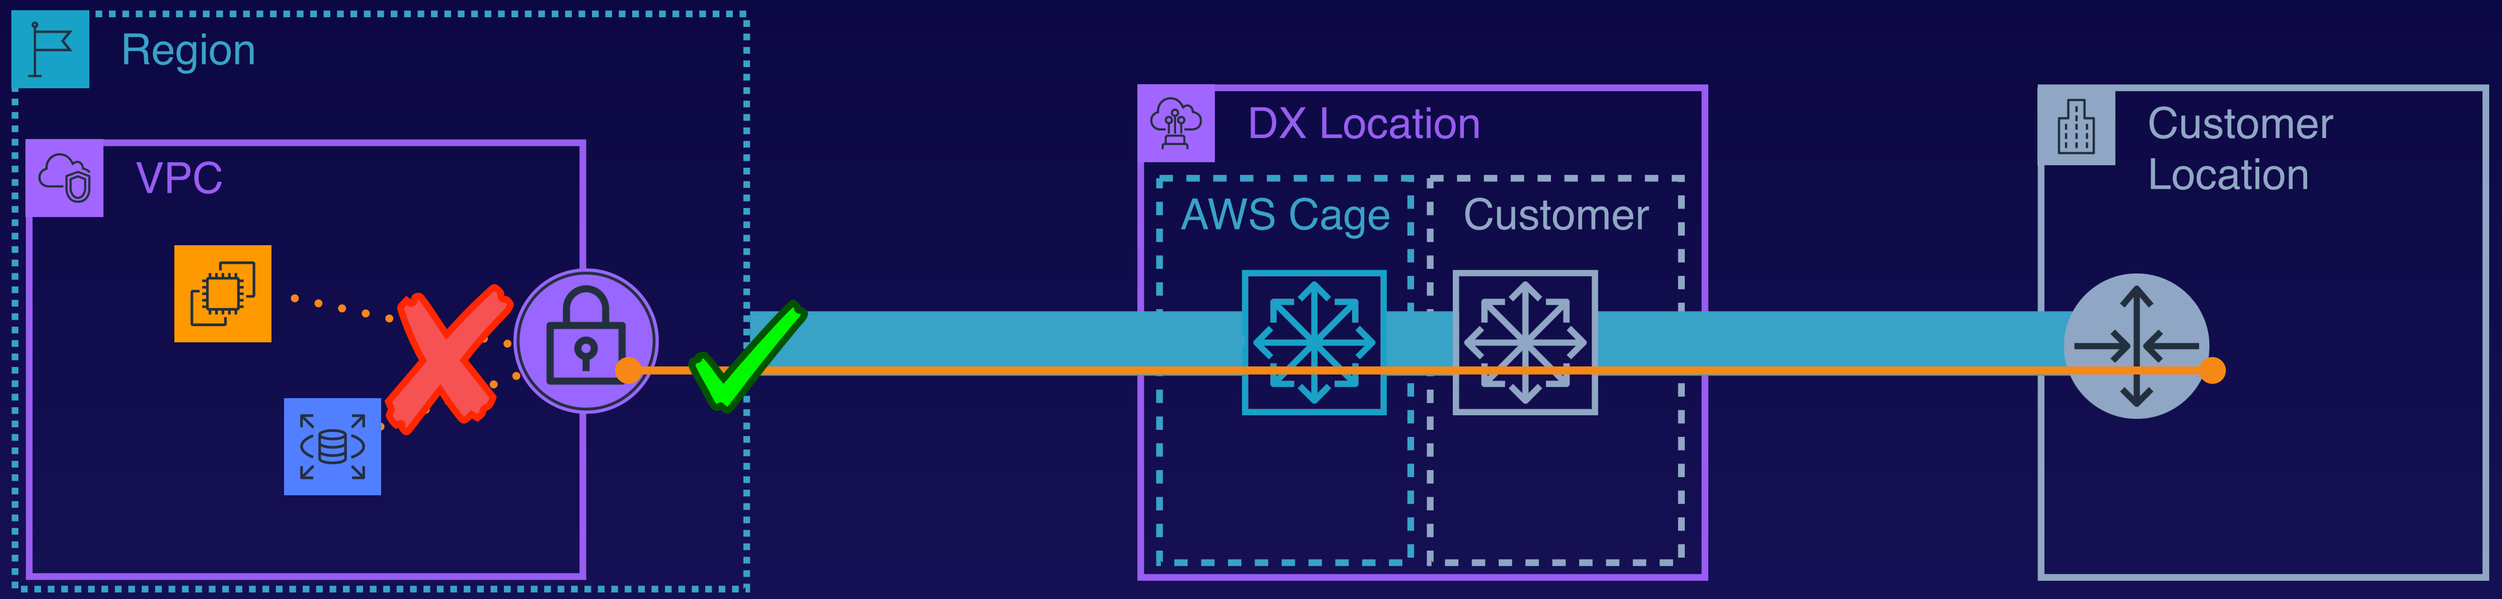

- Connection, VIF, and BGP session are up but you can't reach AWS resources?
  - Routing to AWS resources correctly configured?
  - NACLs and security groups correctly configured?
  - Software at AWS resources correctly configured?

Fast Takeaways

Customer is responsible for implementing architecture to support DX reliability.

LAGs and jumbo frames can be used to improve performance.

AWS does not encrypt DX traffic.

DX offers lower DTO rates but includes hourly port charges.

Customer is responsible for monitoring DX connection health.