Exercise 3 — Attack Chain Reconstruction: Answer

The Narrative, Stage by Stage

Stage 1

"A remote employee connects to the office VPN from a hotel."

Attack technique: This is the setup stage — no attack yet, but the threat surface is already open. The employee is on an untrusted network (hotel Wi-Fi) with no guarantee the network itself is clean.

Residual risk to name: Split tunnelling — if enabled, the employee's traffic to non-corporate sites bypasses the VPN entirely, and a compromised hotel network could intercept that traffic or pivot to the device.

Control that should be in place:

- Device compliance check / NAC before VPN admission (is the endpoint patched? Is EDR running?)
- MFA required at VPN login — so a stolen password alone is not enough

---

Stage 2

"They notice their browser shows a certificate warning for the company intranet — they click through."

Attack technique: Man-in-the-Middle (MitM) / TLS downgrade or SSL stripping. The attacker — positioned on the hotel network — is intercepting the TLS session and presenting a fraudulent certificate. The browser warning is the only signal the employee receives, and they ignore it.

Defensive layer that failed: Layer 1 (Perimeter) and the VPN itself. If the VPN tunnel was properly established *before* the intranet connection, traffic should have been inside the encrypted tunnel and this interception shouldn't have been possible — which suggests either:

- The VPN wasn't fully established before the browser session opened, or
- Split tunnelling allowed the intranet traffic to travel outside the VPN

Controls that should have stopped it:

- **HSTS (HTTP Strict Transport Security)** configured on the intranet server — this prevents the browser from accepting a downgraded or invalid certificate, even if the user clicks through
- **Certificate pinning** for internal applications
- **User awareness training** — clicking through a certificate warning is a critical failure that training should address
- A **zero-trust application proxy** (ZTNA) would have meant the user never connects directly to the intranet at all — they authenticate to the ZTNA broker, which validates both the user *and* the device before brokering application access, eliminating this attack vector entirely

---

Stage 3

"Within 30 minutes, their credentials are captured."

Attack technique: The MitM position allows the attacker to harvest the credentials as they are submitted through the intercepted session — effectively **credential harvesting via session hijacking**. Even if the password is never typed again, the attacker may also capture a valid session token, allowing them to replay it.

Defensive layer that failed: Authentication design. A password alone, even over what the user *thought* was a secure connection, is insufficient.

Controls that should have stopped it:

- **MFA** — even with the password captured, the attacker cannot authenticate without the second factor (TOTP, push notification, hardware token)
- **Mutual TLS (mTLS)** — requires the *client* to also present a certificate, so the server can verify the device, not just the user. A rogue attacker without the client certificate cannot complete the handshake
- **Phishing-resistant MFA** (e.g., FIDO2/WebAuthn hardware keys) is the gold standard — these are cryptographically bound to the legitimate origin domain, so credentials captured via MitM cannot be replayed to the real site

---

Stage 4

"The attacker uses those credentials to log in from a different country."

Attack technique: Credential stuffing / account takeover — the attacker authenticates using the stolen credentials from a completely different geographic location.

Defensive layer that failed: Layer 2 (Network Security) monitoring and contextual access controls.

Controls that should have stopped it:

- **Behavioural analytics (UEBA — User and Entity Behaviour Analytics)** — a login from a country the user has never accessed from, within minutes of a session in the hotel, is a clear anomaly (impossible travel). A UEBA system should flag or block this automatically
- **Conditional Access Policies** — modern identity platforms (tied to ZTNA or IAM) can enforce rules such as: *"If login location deviates from baseline, step up to additional verification or deny"*
- **SIEM correlation** — even without automated blocking, a SIEM correlating the two login events (hotel IP → foreign IP within 30 minutes) should trigger a security alert for the SOC

This is also a key point to raise: **this is exactly the kind of event that a traditional VPN with no post-authentication monitoring would completely miss** — the attacker is now "inside" with valid credentials and the perimeter firewall sees nothing wrong.

Stage 5

"They move laterally to a file server in a different VLAN."

Attack technique: Lateral movement — the attacker, now authenticated, explores the network and pivots from their initial access point to other systems. The fact that they can reach a file server in a *different VLAN* is the critical failure here.

Defensive layer that failed: This is the "chewy on the inside" problem from the Layered Defense slide. The internal network lacked sufficient segmentation enforcement — inter-VLAN routing was either too permissive or unmonitored.

Controls that should have stopped it:

- **Microsegmentation** — even within VLANs, east-west traffic between hosts should require explicit policy permission. A compromised user account should not automatically inherit access to file servers in other segments

- **Least privilege access control** — the user's account should only have access to the specific resources their role requires. If the compromised account had no business reason to access that file server, the connection should have been denied
- **Internal IDS/IPS** — an IDS monitoring *internal* east-west traffic (not just north-south at the perimeter) should detect anomalous access patterns — a user suddenly scanning or accessing multiple file shares is a strong indicator of compromise
- **802.1X and NAC** — verifying device identity and compliance at the point of network access, not just at the perimeter, would have challenged the attacker's ability to move between segments

---

Stage 6

"Three hours later, a database server begins sending large volumes of data to an external IP."

Attack technique: Data exfiltration — the attacker has reached a high-value target (the database server) and is extracting data outbound. The three-hour gap suggests they spent time identifying and staging the data before exfiltrating.

Defensive layer that failed: Layer 5 (Data Security) and outbound monitoring. The data left the network undetected for long enough to be a significant breach.

Controls that should have stopped it:

- **Data Loss Prevention (DLP)** — should inspect outbound traffic for sensitive data patterns (PII, financial records, bulk data transfers) and block or alert on policy violations
- **Egress filtering on the firewall** — database servers should virtually never initiate outbound connections to arbitrary external IPs. A firewall rule denying outbound traffic from the DB server subnet to the internet (except approved destinations) would have blocked this entirely. This is a configuration failure — **servers that only need to receive connections should not be permitted to initiate them**
- **SIEM alert on data volume anomaly** — a database server suddenly transferring gigabytes outbound at 2am is an obvious anomaly that a SIEM rule should catch
- **Network behaviour analysis** — tools watching for unusual flow patterns (a server that typically sends 10MB/day suddenly sending 10GB) would flag this

---

Full Chain Summary Table

Stage	Attack Technique	Layer That Failed	Key Control That Would Have Stopped It
1	Untrusted network exposure	Pre-admission	NAC / device compliance check before VPN
2	MitM / TLS interception	Perimeter + VPN design	HSTS, mTLS, ZTNA broker — or properly enforced full-tunnel VPN
3	Credential harvesting	Authentication design	MFA (ideally FIDO2/phishing-resistant)
4	Account takeover from foreign IP	Identity monitoring	UEBA, conditional access, SIEM impossible-travel alert
5	Lateral movement across VLANs	Internal segmentation	Microsegmentation, least privilege, internal IDS
6	Data exfiltration	Data security + egress	Internet blocking from D’base server, DLP, egress filtering, SIEM volume anomaly alert

---

The Overarching Lesson to Land

The entire chain succeeds because each layer *individually* was either absent, misconfigured, or produced an alert that no one acted on. This is the core argument for defence in depth — **no single control stops a determined attacker, but multiple overlapping controls make each step progressively harder and louder**. The attacker had six opportunities to be stopped or detected before a single byte of data left the building. The fact that they weren't is a failure of architecture, not just technology.

This also makes the case for **ZTNA over traditional VPN** cleanly: under a zero-trust model, stages 2, 4, and 5 are dramatically harder — the user never gets "on the network," lateral movement has no network to traverse, and every application access is re-verified continuously rather than trusted implicitly after the initial handshake.