Redacted Network & Infrastructure Penetration Test Report (Sample)

Contents

Attestation: This document is a sanitized, representative sample. All organization names, domains, IPs, credentials, and timestamps have been replaced or removed. Evidence excerpts may be cropped or blurred.

1 Executive Summary

This assessment simulated adversary activity against **Acme Manufacturing, Inc.** (placeholder entity) with a focus on **external perimeter, identity/remote access, and core network controls**. Overall risk is **High** due to an exploitable edge appliance, weak legacy authentication paths, and gaps in network hardening.

1.1 Highlights

- Externally exposed **VPN appliance** fingerprinted to a version family with a **critical RCE** (patch available).
- **NTLM relay** feasible on internal paths due to missing SMB signing and misconfigured LDAP protections, enabling **credential reuse and lateral movement**.
- **Unauthenticated management services** detected on select hosts, increasing compromise and persistence opportunities.
- **Weak TLS** parameters and **legacy protocols** of record (TLS 1.0/1.1) on a subset of endpoints.
- Cloud perimeter review found **public object storage list/read** on a non-sensitive bucket (still a governance issue).

1.2 Business impact (summary)

Loss of service, data disclosure, and potential domain compromise stemming from identity abuse and perimeter exploitation. Remediation should prioritize **edge patching/MFA**, **identity hardening (SMB/LDAP/TLS)**, and **segmentation**.

2 Technical Summary

2.1 Scope

- External perimeter: selected IP ranges and FQDNs
- Remote access/identity: VPN gateway, RDP/SSH jump paths, IdP sign-in flows (black-box)
- Core network controls: SMB/LDAP exposure, management interfaces (safe checks only)
- Limited cloud perimeter sampling for storage and logging controls

2.2 Caveats

- DoS/destructive testing excluded; only safe exploitation demonstrated.
- Credential brute-forcing excluded; proof paths used controlled test accounts/tokens where required.
- Some internal checks performed via a controlled test vantage point emulating a foothold.
- Evidence included here is **redacted**.

Post-Assessment Cleanup Temporary test accounts/tokens were revoked; artifacts purged from test hosts; indicators provided for SOC tuning.

3 Table of Findings

ID	Title	Severity	Likelihood	Surface	Status
F-001	VPN Appliance Exposed with Critical RCE	**Critical**	High	External perimeter	Unfixed
F-002	NTLM Relay Feasible	**High**	High	Internal lateral paths	Unfixed
F-003	Unauthenticated Management Service Exposure	**High**	Medium	External & internal segments	Unfixed
F-004	Weak TLS Configuration	**Medium**	Medium	External HTTPS endpoints	Unfixed
F-005	Public Read on Cloud Object Storage	**Low**	Medium	Cloud perimeter	Unfixed

4 Risk Ratings

We use a composite risk model and also calculate a CVSS base score for each finding.
Primary score: `Risk = Severity × Exploitability × Exposure`
Reference score: `CVSS (v3.1 by default; v4.0 on request)`

Final categorical rating (**Critical / High / Medium / Low / Informational**) is derived from the **primary score**, then **harmonized** with CVSS (see §4.4).

4.1 Dimensions (what we measure)

Severity (Business Impact) — Consequence to confidentiality, integrity, availability, safety, and regulatory posture if exploited.

- **5 – Catastrophic:** Domain/tenant compromise, widespread outage, regulated data breach, existential business risk.
- **4 – Major:** Material data exposure, privileged escalation, sustained service degradation.
- **3 – Moderate:** Limited data exposure, contained lateral movement, disruption with workarounds.
- **2 – Minor:** Small blast radius, low-sensitivity data, localized disruption.
- **1 – Minimal:** Negligible business effect.

Exploitability — Effort, prerequisites, and reliability of attack.

- **5 – Trivial:** Remote, unauthenticated; public exploit; low skill/time.
- **4 – Easy:** Minimal prerequisites; partial auth or common misconfig.
- **3 – Practical:** Needs a user-in-path, single misstep, or limited foothold.
- **2 – Difficult:** Multiple conditions, timing windows, scarce tooling.

- **1 – Improbable:** Theoretical or highly constrained.

Exposure — Who can reach the vulnerable surface.

- **5 – Internet-Facing (broad)**
- **4 – Partner/External**
- **3 – Internal (wide)**
- **2 – Internal (restricted)**
- **1 – Lab/Non-Prod**

Optional Modifiers (documented, not multiplied): Compensating Controls, Detectability, and Evidence Confidence can justify nudging a rating up/down one tier in edge cases.

CVSS Base Score (Reference Dimension) — Calculated per **CVSS v3.1** (AV/AC/PR/UI/S/C/I/A) unless you request v4.0. We publish the **vector** and **base score** for each finding to aid cross-vendor comparisons and ticketing systems.

4.2 Scoring & Category Thresholds

Primary product score: `Score = S` $\times$ `E` $\times$ `X` (range 1–125)

Product Score	Category	Typical Response Target*
100–125	**Critical**	Contain/patch **0–7 days**
60–99	**High**	Remediate $\leq$**30 days**
20–59	**Medium**	Remediate $\leq$**90 days**
5–19	**Low**	Backlog / hardening wave
1–4	**Info**	Track only

*Targets are defaults; SLAs may supersede.

Why product, not sum? Multiplication elevates truly dangerous combinations (e.g., internet-facing $\times$ trivial exploit $\times$ major impact) and prevents low-exposure/low-impact issues from inflating unduly.

4.3 CVSS Mapping (reference bands)

Unless otherwise noted, we use **CVSS v3.1** severity bands:

CVSS Base Score	CVSS Severity
9.0–10.0	**Critical**
7.0–8.9	**High**
4.0–6.9	**Medium**
0.1–3.9	**Low**
0.0	**None**

v4.0 option: If requested, we will compute **CVSS v4.0** (which refines environmental/contextual scoring). Bands are materially similar for base severity; we'll note any differences.

4.4 Harmonization Rule (Primary × CVSS)

1. Derive the **Primary Category** from the product score table (§4.2).
2. Derive the **CVSS Category** from the CVSS bands (§4.3).
3. **Final Category = max(Primary, CVSS)** by severity tier, **unless** clearly documented Compensating Controls/Detectability justify at most **one-tier reduction** (e.g., strong segmentation + airtight monitoring). All exceptions are explicitly noted.

This approach preserves operational realism (exposure and exploitability in your environment) while remaining interoperable with industry standards and ticketing/IR pipelines.

4.5 Examples (this engagement's pattern)

- **Edge RCE on VPN (unauth'd, internet):** `S=5, E=5, X=5 → 125 (Primary: Critical)`; **CVSS v3.1** example vector *AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:H/A:H* → **9.8 (Critical)** → **Final: Critical**.
- **NTLM Relay (internal wide, practical):** `S=4, E=3–4, X=3 → 36–48 (Primary: High)`; **CVSS** example *AV:A/AC:L/PR:N/UI:N/S:U/C:H/I:H/A:L* → **High** → **Final: High**.
- **Weak TLS (internet, compliance risk):** `S=2, E=3, X=4 → 24 (Primary: Medium)`; **CVSS** example *AV:N/AC:H/PR:N/UI:N/S:U/C:L/I:N/A:N* → **Low–Medium** depending on context → **Final: Medium** (unless compensating controls warrant Low).

  Note: Example vectors are illustrative; each finding in §5 lists its actual vector and both scores.

4.6 4.6 Status, Confidence, and Residual Risk

Each finding includes:

- **Status:** *Unfixed / In progress / Fixed – Pending Retest / Verified Fixed*
- **Evidence Confidence:** *High / Medium / Low*
- **Residual Risk:** Updated **Primary** and **CVSS** after partial fixes or new controls.

4.7 4.7 Re-rating Triggers

We re-rate when any of the following change: exposure surface (e.g., moved off internet), exploitability (e.g., MFA, SMB signing), severity (e.g., sensitive data confirmed/removed), or CVSS vector components.

5 Finding Details

Scoring legend: Primary = `S` × `E` × `X` (Severity, Exploitability, Exposure). CVSS = v3.1 base (vector shown). **Final Category = max(Primary, CVSS)** unless a one-tier reduction is explicitly justified by strong compensating controls.

5.1 F-001: VPN Appliance Exposed with Critical RCE (Version Family Vulnerable) — Critical

Scoring

- **Primary:** `S=5 × E=5 × X=5 = 125` → **Critical**
- **CVSS v3.1: 9.8 (Critical)** — `AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:H/A:H`
- **Final Category: Critical**
- **Status:** Unfixed **Evidence Confidence:** High

Description — The external VPN gateway at `vpn.acme.example` presents responses and banners consistent with a vendor build affected by a widely exploited **remote code execution** flaw. Proof-of-concept exploitation was **not** executed per rules of engagement; risk remains material pending patching.
Impact — Edge compromise with potential for credential interception and lateral movement.
Evidence

```
# Request/response excerpts (sanitized)
Host: vpn.acme.example
GET /status HTTP/1.1
<headers redacted>

# Fingerprints
Banner: <redacted vendor/version family>
Favicon-Hash: <redacted>
Status-Metadata: build=<redacted>; branch=<vuln-family>
```

Remediation — Patch/upgrade to vendor fixed release; enforce **MFA** for all remote access; disable legacy portals; collect telemetry (WAF/VPN logs) to SIEM; add geo/behavioral access controls.
Mappings — ATT&CK T1190 (Exploit Public-Facing App); NIST PR.IP-12.

5.2 F-002: NTLM Relay Feasible (SMB Signing Disabled + LDAP Protections Missing) — High

Scoring

- **Primary:** `S=4 × E=4 × X=3 = 48` → **High**
- **CVSS v3.1: 8.3 (High)** — `AV:A/AC:L/PR:N/UI:N/S:U/C:H/I:H/A:L`
- **Final Category: High**
- **Status:** Unfixed **Evidence Confidence:** High

Description — Internal scan and handshake review indicated **SMB signing disabled** on multiple hosts and no enforced channel bindings on directory services; combined, this enables **NTLM relay** under realistic preconditions.
Impact — Privilege escalation and lateral movement; potential to obtain useful tokens or write operations in directory-adjacent services.
Evidence

```
# SMB negotiation (sanitized)
SMB.Signing: Disabled
SMB.Version: 3.x (varies)

# Directory service checks
LDAP.SimpleBind: Allowed
LDAP.ChannelBinding: Not Enforced
LDAP.Signing: Not Required

# Discovery notes
LLMNR/NetBIOS: responder-style discovery suppressed during test window
```

Remediation — Enforce SMB signing (domain/host GPO); require LDAP channel binding/signing; phase out NTLMv1; prefer Kerberos; monitor for LLMNR/NetBIOS abuse; enable Protected Users/ESAE patterns where feasible.
Mappings — ATT&CK T1557 (Adversary-in-the-Middle), T1078 (Valid Accounts); NIST PR.AC-3.

5.3 F-003: Unauthenticated Management Service Exposure (RDP/SSH subsets) — High

Scoring

- **Primary:** `S=4 × E=3 × X=4 = 48` → **High**
- **CVSS v3.1: 8.0 (High)** — `AV:N/AC:L/PR:N/UI:N/S:U/C:L/I:L/A:L`
- **Final Category: High**
- **Status:** Unfixed **Evidence Confidence:** Medium

Description — Select hosts exposed management services (**RDP**, **SSH**) to broad networks without strong broker/allow-list controls; banner checks indicated default or weak guardrails.
Impact — Increased likelihood of brute-force, credential stuffing, or exploitation of protocol/device vulnerabilities; facilitates persistent footholds.
Evidence

```
# nmap (shallow, sanitized)
3389/tcp open  ms-wbt-server  syn-ack
22/tcp   open  ssh            syn-ack

# Hardening flags observed
RDP: Network Level Authentication (NLA) = Not Enforced (subset)
SSH: Default banner present; host key rotation policy unclear
Access Controls: No bastion/broker required on subset of hosts
```

Remediation — Place management behind bastion/Broker; enforce MFA; restrict by source; enable network-level auth; rotate host keys; harden golden images.
Mappings — ATT&CK T1021 (Remote Services); NIST PR.AC-5.

5.4 F-004: Weak TLS Configuration (Legacy Protocols/Ciphers) — Medium

Scoring

- **Primary:** `S=2 × E=3 × X=4 = 24` → **Medium**
- **CVSS v3.1: 4.0 (Medium)** — `AV:N/AC:H/PR:N/UI:N/S:U/C:L/I:N/A:N`
- **Final Category: Medium**
- **Status:** Unfixed **Evidence Confidence:** Medium

Description — Several HTTPS endpoints allow **TLS 1.0/1.1** and legacy ciphers, widening downgrade surface and compliance risk.
Impact — Reduced confidentiality guarantees; audit/compliance findings.
Evidence

```
# Scan excerpt (sanitized)
Protocols:
  TLSv1:   ENABLED
  TLSv1.1: ENABLED
  TLSv1.2: ENABLED
Ciphers (subset):
  TLS_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA  (weak)
  TLS_RSA_WITH_AES_128_CBC_SHA         (legacy)
HSTS: Not enforced on subset
```

Remediation — Enforce TLS 1.2+ (prefer 1.3); remove weak ciphers; enable **HSTS**; standardize server profiles.
Mappings — NIST SP 800-52r2 (TLS Baselines); OWASP ASVS v4.0.3 (V9: Communications — V9.1, V9.2).

5.5 F-005: Public Read on Cloud Object Storage (Governance Gap) — Medium

Scoring

- **Primary:** `S=2 × E=3 × X=4 = 24` → **Medium**
- **CVSS v3.1: 5.3 (Medium)** — `AV:N/AC:L/PR:N/UI:N/S:U/C:L/I:N/A:N`
- **Final Category: Medium**
- **Status:** Unfixed **Evidence Confidence:** Medium

Description — A bucket used for build artifacts permits unauthenticated `GetObject` on a `public/` prefix. Files observed were non-sensitive; configuration still introduces governance risk.
Impact — Accidental disclosure if sensitive artifacts are later uploaded or tagging drifts.
Evidence

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Principal": "*",
    "Action": "s3:GetObject",
    "Resource": "arn:aws:s3:::acme-artifacts/public/*"
  }]
}
```

Remediation — Block public access; require signed URLs or IAM principals; enforce preventive controls (SCP/Config); continuous checks.
Mappings — ATT&CK T1530 (Data from Cloud Storage); CIS Benchmarks.

6 Remediation Roadmap (Prioritized)

Principles: (a) **Exploit-stopper first** (edge & identity), (b) **Reduce blast radius** (brokered admin + segmentation), (c) **Raise the floor** (crypto baselines & guardrails), (d) **Continuously verify** (telemetry + retest).
Discipline: Capture exceptions with **owner, reason, scope, expiry date**. Pilot changes in **rings** (pilot → staged → broad).

6.1 Patch/Upgrade VPN Edge; Enforce MFA; Tighten Remote Access *(Critical)*

Owners: Network Eng • IAM • SecOps
Pre-checks: Inventory gateways; confirm HA/rollback path; export configs; confirm vendor fixed versions.
Implementation

- Upgrade/patch to the vendor's fixed release; remove legacy web portals/modules.
- Enforce **MFA** on all remote access (users and admins), including VPN client and any web SSO to the gateway.
- Restrict sources to known geo/ASN or partner IPs; consider client certificates for admins.

- Disable local accounts where SSO exists; standardize strong TLS profiles on the gateway.
- Shorten idle/absolute session lifetimes; disable split tunneling unless formally excepted.

Validation

1. Version/build matches fixed release.
2. Legacy endpoints return 403/404 (blocked).
3. Login without MFA is blocked.
4. Only approved TLS versions/ciphers are offered.
5. Split tunneling disabled (or documented exception).
6. SIEM receiving auth/admin logs; synthetic failed-MFA alert fires.

Telemetry

- VPN auth/admin events to SIEM; alerts for impossible travel, repeated failures, config drift.
- Dashboard: VPN auths by source, MFA failures, admin changes.

6.2 Eliminate NTLM Relay; Harden Directory Paths *(High)*

Owners: AD/Windows Eng • IAM
Pre-checks: Identify systems with SMB signing off; inventory LDAP channel binding/signing; ensure SOC can monitor LLMNR/NetBIOS.

Implementation

- Enable **SMB signing** via GPO (required on Tier-0; enabled elsewhere).
- Enforce **LDAP channel binding** and **LDAP signing**; prefer **LDAPS** for apps.
- Reduce/disable **NTLM** (block NTLMv1; prefer Kerberos).
- Disable **LLMNR/NetBIOS** where feasible; harden WPAD.
- Tiering: protect admin accounts (Protected Users, PAW/bastion), isolate Tier-0 services.

Validation

1. SMB negotiation shows Signing: Required (servers) / Enabled (clients).
2. LDAP simple binds without signing are blocked.
3. NTLMv1 attempts are denied; NTLMv2 events trend down.
4. Responder/LLMNR tests capture/relay attempts → none.
5. LDAPS enforced for directory-bound applications (sample verified).
6. RSoP/GPO export attached for Tier-0 scope.

Telemetry

- Alerts: unsigned SMB sessions, LDAP simple binds, NTLM usage spikes; track downward trend post-change.

6.3 Broker & Restrict Management Access (RDP/SSH) *(High)*

Owners: Platform Eng • Identity • Network
Pre-checks: Enumerate RDP/SSH exposures; define admin groups; select bastion/broker (e.g., RD Gateway/Privileged Access Workstation, or SSH bastion).

Implementation

- Require all admin access through the **broker** with MFA; block direct management from untrusted subnets.
- Firewall allow-lists for broker egress only; deny broad inbound management to servers.
- **RDP**: enforce NLA, TLS, clipboard/drive controls; record privileged sessions where permitted.
- **SSH**: disable password auth; require keys or **short-lived certs**; rotate host keys; standardize banners.
- Bake hardening into **golden images**; auto-join to broker policy.

Validation

1. Direct RDP/SSH from workstations is blocked.
2. RDP via broker with MFA is permitted and logged.
3. SSH shows `PasswordAuthentication=no`, `PubkeyAuthentication=yes`.
4. Sample servers reflect required policies (NLA on; banners updated).
5. PAM/PIM just-in-time elevation demonstrated; privileges expire automatically.
6. Session recording sample reviewed (privacy/compliance banner present).

Telemetry

- Broker access logs + session recordings (where allowed); alerts on direct management attempts; firewall/broker policy change audit.

6.4 Harden TLS to 1.2+/1.3; Remove Weak Ciphers; Add HSTS *(Medium)*

Owners: App/Platform Eng
Pre-checks: Inventory endpoints; identify legacy clients; plan blue/green where possible.

Implementation

- Disable TLS 1.0/1.1; prefer TLS 1.3 where supported.
- Remove RC4/3DES/NULL/EXPORT; prefer ECDHE + AEAD suites.
- Enable **HSTS** with conservative `max-age`; expand to `includeSubDomains` after verification.
- Centralize TLS profiles (ingress/reverse proxy) to reduce config drift.
- Add **OCSP stapling** and certificate linting/expiry alerts.

Examples

```
ssl_protocols TLSv1.2 TLSv1.3;
ssl_ciphers HIGH:!aNULL:!MD5:!3DES;
add_header Strict-Transport-Security "max-age=15552000; includeSubDomains" always;
```

Validation

1. TLS scans show only TLS 1.2/1.3.

2. No deprecated ciphers accepted.
3. HSTS present on HTTPS responses (pilot → full rollout).
4. Cert expiry alerts active; OCSP stapling validated on sample endpoints.

Telemetry

- Continuous TLS scans; alert on regressions/drift; monitor failed handshakes by client type.

6.5 Tighten Cloud Object Storage Policy; Continuous Governance *(Low)*

Owners: Cloud Eng • SecOps
Pre-checks: Identify buckets with public ACLs/policies; classify contents; define exceptions (CDN/public assets).

Implementation

- Enable account- and bucket-level **Block Public Access**.
- Replace `Principal="*"` with role-based policies; require **signed URLs** for distribution.
- Preventive controls: **SCP/Org Policies** to block public ACLs; deploy **Config/Policy** rules for drift detection.
- Centralize server-access logs to a private bucket; enable default encryption and lifecycle for logs.

Examples

```
{
  "BlockPublicAcls": true,
  "IgnorePublicAcls": true,
  "BlockPublicPolicy": true,
  "RestrictPublicBuckets": true
}
```

Validation

1. Public GET to former object paths is denied.
2. Access via signed URL/role is permitted and logged.
3. Config/Policy rule findings = 0.
4. Org-level SCP prevents creating public buckets.
5. Default encryption + lifecycle policies show **Enabled**.

Telemetry

- CloudTrail + Access Logs; alerts on `PutBucketPolicy` that introduces public access; weekly compliance report.

6.6 Cross-Cutting Actions

- **Change Control & Rollback:** document back-out plans; stage in non-prod; promote by rings.
- **SLA Mapping:** Critical $\leq$7 days • High $\leq$30 • Medium $\leq$90 • Low in hardening wave.
- **KPIs:** % endpoints at TLS baseline; % servers with SMB signing required; broker adoption rate; NTLM usage trend; public bucket count = 0.
- **Retest Windows:** schedule verification scans/tests after each change set; update **Residual Risk** in §5.
- **Documentation:** commit hardened configs/playbooks; attach change tickets; record exceptions with expiry dates.

7 Contact Info

Yugen Risk Advisors - Security: security@yugenrisk.com - Primary: +1 (239) 427-1486 - Web: https://yugenrisk.com