CONFIDENTIAL — DISTRIBUTION RESTRICTED

PENETRATION TEST REPORT

Security Assessment of 192.168.1.12

DOCUMENT ID
NPR-20260601-E5A7FF

DATE ISSUED
June 1, 2026

ENGAGEMENT TYPE
Autonomous Network Penetration Test (Tier 2)

OVERALL RISK RATING
CRITICAL

PREPARED FOR
- Customer Organisation -

PREPARED BY
NeuraPent Autonomous Security

REPORT VERSION
1.0

NEURAPENT SECURITY PLATFORM

PAGE 1 OF N

Document Control

VERSION	DATE	CHANGE DESCRIPTION	AUTHOR
1.0	01/06/2026	Initial release of penetration test findings	NeuraPent Engine

DISTRIBUTION & HANDLING

This report contains sensitive security information including evidence of exploited vulnerabilities. It is restricted to authorised personnel within the customer organisation. Reproduction, redistribution or disclosure outside the explicit recipient list requires written approval from the engagement sponsor.

Table of Contents

1. Executive Summary

NeuraPent conducted an autonomous penetration test against **192.168.1.12** commencing on May 31, 2026. The objective of the engagement was to identify exploitable security weaknesses and demonstrate the impact of those weaknesses through controlled exploitation, in line with the Penetration Testing Execution Standard (PTES).

The assessment was completed in 35 engagement turns, covering 25 exposed services. A total of **24 findings** were identified, of which **4 are rated Critical** and **6 are rated High**. No interactive shells were obtained during the engagement.

2 multi-step attack chains were identified and verified, demonstrating realistic exploitation paths from initial access to material compromise. These narratives are documented in Section 7.

The overall risk to the in-scope environment is rated **Critical**. Immediate remediation is required for all Critical and High findings. Detailed technical findings are presented in Section 6.

KEY FINDINGS

- **[CRITICAL]** Apache Server Configuration Disclosure — 192.168.1.12:80/http
- **[CRITICAL]** Unauthenticated Redis Access — 192.168.1.12:6379/redis
- **[CRITICAL]** Unauthenticated Elasticsearch API Access — 192.168.1.12:9200/elasticsearch
- **[CRITICAL]** Attack Chain: Unauthenticated Redis to RCE and Internal Pivot — 192.168.1.12:6379/redis
- **[HIGH]** Writable Network Share — 192.168.1.12:139/?
- **[HIGH]** CVE-2023-38709 — 192.168.1.12:80/http

2. Scope & Objectives

The engagement was scoped to the assets listed below. No other systems, networks, or applications were tested. Engagement Tier 2 was authorised, permitting active verification and proof-of-concept exploitation but excluding destructive payloads or denial-of-service techniques.

IN-SCOPE ASSETS

HOST / NETWORK	DISCOVERED SERVICES	ENGAGEMENT WINDOW
192.168.1.12	25 TCP services (80/tcp, 135/tcp, 139/tcp, 443/tcp, 445/tcp, 623/tcp, 902/tcp, 912/tcp, 2179/tcp, 3306/tcp, 5000/tcp, 5040/tcp, 5440/tcp, 6379/tcp, 7070/tcp, 7474/tcp, 7687/tcp, 9200/tcp,	May 31, 2026

16992/tcp, 49664/tcp, 49665/tcp, 49666/tcp, 49667/tcp, 49668/tcp, 49676/tcp)

ENGAGEMENT OBJECTIVES

- Identify externally exposed services and software versions on the target
- Identify and validate exploitable security weaknesses
- Demonstrate realistic business impact through controlled exploitation
- Provide actionable remediation guidance prioritised by risk

3. Methodology

Testing was performed using the NeuraPent agentic assessment framework, which follows the Penetration Testing Execution Standard (PTES) phased model. The framework operates an autonomous AI orchestrator that selects, executes, and iterates on industry-standard offensive security tools within a sandboxed Kali Linux environment. Each phase's output informs the next, allowing the engagement to dynamically adapt to discovered conditions in the same way a human red-team operator would.

#	PHASE	DESCRIPTION
01	**Reconnaissance**	Map the attack surface
02	**Vulnerability Audit**	Identify weaknesses
03	**AI Classifier**	Prioritize & cluster
04	**Verify & Exploit**	Confirm exploitability
05	**Post-Exploitation**	Extract loot & shells
06	**Forensic Cleanup**	Remove artifacts
07	**Report Generation**	Compile evidence

Tools employed during this engagement include, but are not limited to: *Nmap, Nikto, Searchsploit, SQLMap, Metasploit Framework, Hydra, Gobuster, custom verification scripts and protocol-specific clients.* All exploitation activity was conducted within the boundaries of the authorised engagement tier.

4. Risk Rating Methodology

Findings are rated according to the Common Vulnerability Scoring System (CVSS) v3.1 base score and the resulting business risk in the customer's environment. The mapping is as follows:

RATING	CVSS RANGE	DEFINITION & RECOMMENDED ACTION
CRITICAL	9.0 – 10.0	Direct compromise of confidentiality, integrity or availability without prior access. **Immediate remediation required (within 7 days).**
HIGH	7.0 – 8.9	Significant impact, typically requiring limited prior access or user interaction. **Remediate within 30 days.**
MEDIUM	4.0 – 6.9	Moderate impact; usually a building block towards a higher-severity issue. Remediate within next maintenance window.
LOW	0.1 – 3.9	Minor information disclosure or hardening shortfall. Address as part of routine hardening.
INFO	0	Informational; no direct security impact, included for completeness.

5. Summary of Findings

The following table lists all 24 findings identified during the engagement, ordered by severity. Each finding is documented in detail in Section 6.

ID	SEVERITY	FINDING	CVSS	STATUS
F-001	CRITICAL	Apache Server Configuration Disclosure (http/80)	9.1	Potential
F-002	CRITICAL	Unauthenticated Redis Access (redis/6379)	9.8	Potential
F-003	CRITICAL	Unauthenticated Elasticsearch API Access (elasticsearch/9200)	9.8	Potential
F-004	CRITICAL	Attack Chain: Unauthenticated Redis to RCE and Internal Pivot (redis/6379)	9.8	Potential
F-005	HIGH	Writable Network Share (?/139)	–	Potential
F-006	HIGH	CVE-2023-38709 (http/80)	7.3	Potential
F-007	HIGH	CVE-2025-49812 (ssl/http/443)	7.4	Potential
F-008	HIGH	CORS Misconfiguration Allows Credentialed Cross-Origin Data Theft (http/5000)	6.9	Potential
F-009	HIGH	Intel Active Management Technology (AMT) Vulnerable to Privilege Escalation (unknown/623)	8.2	Potential
F-010	HIGH	Writable Network Share (multi-service/139)	9.8	Potential
F-011	MEDIUM	Weak TLS Cipher Suites Supported (Logjam) (http/443)	3.0	Potential
F-012	MEDIUM	Sensitive API Endpoints and Service Details Disclosed on Root Endpoint (http/5000)	6.5	Potential
F-013	MEDIUM	Intel Active Management Technology (AMT) Version Disclosure (http/16992)	4.2	Potential
F-014	MEDIUM	VMware Authentication Daemon Version Disclosure Indicates End-of-Life Software (ssl/vmware-auth/902)	3.0	Potential
F-015	MEDIUM	VMware Authentication Daemon Version 1.0 Information Disclosure (vmware-auth/912)	4.2	Potential
F-016	MEDIUM	Outdated MariaDB Version Detected (mysql/3306)	3.0	Potential
F-017	MEDIUM	AnyDesk Remote Desktop Service Identified (ssl/realserver?/7070)	3.3	Potential
F-018	MEDIUM	Neo4j Service Leaks Version and Internal IP Address (unknown/7474)	6.5	Potential
F-019	MEDIUM	Neo4j Service Configured with Insecure CORS Policy (unknown/7474)	3.8	Potential
F-020	MEDIUM	Insecure CORS Policy Allows Cross-Domain Data Access (bolt?/7687)	3.8	Potential
F-021	MEDIUM	Elasticsearch Version and Cluster Information Disclosure (elasticsearch/9200)	3.3	Potential
F-022	LOW	Verbose Server and Technology Banners Disclosed (http/80)	–	Potential
F-023	LOW	Missing Subresource Integrity (SRI) (http/80)	–	Potential
F-024	LOW	Verbose HTTP Server Header Discloses Technology (http/5000)	–	Potential

6. Detailed Findings

F-001 Apache Server Configuration Disclosure

SEVERITY	CVSS V3.1	STATUS	ASSET
CRITICAL	9.1	Potential	192.168.1.12:80/http

DESCRIPTION

The Apache `mod_info` module is enabled and publicly accessible at the `/server-info` endpoint. This page discloses the entire Apache server configuration, including loaded modules, virtual host settings, content of configuration files, and internal file system paths.

AFFECTED ASSET

Host	192.168.1.12
Port	80
Service	http (Apache httpd 2.4.58 ((Win64) OpenSSL/3.1.3 PHP/8.2.12))

EVIDENCE

```
$ nuclei -t http/misconfiguration/apache/apache-server-status.yaml -u http://192.168.1.12
[apache-server-status] [http] [low] http://192.168.1.12:80/server-info

$ curl -s http://192.168.1.12/server-info | grep 'Server Root'
<dt><strong>Server Root:</strong></dt> <dd><tt>C:/xampp/apache</tt></dd>
```

BUSINESS IMPACT

An attacker can obtain a complete blueprint of the web server's configuration. This information can be used to identify specific weaknesses, find other vulnerabilities, and plan targeted attacks against the server and its applications. The disclosure of internal paths like `C:/xampp/apache` confirms the OS and installation directory, aiding in path traversal or file inclusion attacks.

REMEDIATION

Disable the `mod_info` handler to prevent configuration disclosure. In the Apache configuration (e.g., `conf/extra/httpd-info.conf`), comment out or remove the `<Location /server-info>` block and restart the Apache service. Example: ```apache #<Location /server-info> # SetHandler server-info # Require local #</Location> ```

F-002 Unauthenticated Redis Access

SEVERITY	CVSS V3.1	STATUS	ASSET
CRITICAL	9.8	Potential	192.168.1.12:6379/redis

DESCRIPTION

The Redis service is configured to allow unauthenticated connections. The service is also bound to all network interfaces, making it accessible to any user on the network. The successful execution of the 'redis-info' script, which retrieves detailed server information, confirms that no password is required to issue commands.

AFFECTED ASSET

Host	192.168.1.12
Port	6379
Service	redis (Redis key-value store 7.4.8)

EVIDENCE

```
$ nmap -p 6379 --script redis-info 192.168.1.12
| redis-info:
|   Version: 7.4.8
|   Operating System: Linux 6.6.114.1-microsoft-standard-WSL2 x86_64
|   Role: master
|   Bind addresses:
|     *
|     -::*
```

BUSINESS IMPACT

An attacker can connect to the Redis database without authentication and gain full control. This allows the attacker to read, modify, and delete all data stored within the database, which may include sensitive information like session tokens, user data, or application secrets. An attacker could also flush the entire database (FLUSHALL), causing a denial of service.

REMEDIATION

1. Enforce authentication by setting a strong password in the Redis configuration file (`redis.conf`). Uncomment and set the `requirepass` directive: `requirepass YOUR_VERY_STRONG_PASSWORD_HERE` 2. Restrict network access. If Redis only needs to be accessed by the local machine, change the `bind` directive in `redis.conf` to listen only on the loopback interface: `bind 127.0.0.1 ::1` If remote access is required, bind to a specific private IP address and use firewall rules to strictly limit connections to trusted source IPs. 3. After making changes, restart the Redis service for them to take effect.

F-003 Unauthenticated Elasticsearch API Access

SEVERITY	CVSS V3.1	STATUS	ASSET
CRITICAL	9.8	Potential	192.168.1.12:9200/elasticsearch

DESCRIPTION

The Elasticsearch REST API on port 9200 is publicly accessible without any authentication. The provided output shows that it was possible to list all indices in the cluster, including one named 'neurapent_loot'. This indicates a complete lack of access control, allowing any user to perform administrative actions, including reading, modifying, and deleting all data stored in the cluster.

AFFECTED ASSET

Host	192.168.1.12
Port	9200
Service	elasticsearch (Elasticsearch REST API 8.12.2 (name: 68930e1807f4; cluster: docker-cluster; Lucene 9.9.2))

EVIDENCE

```
$ curl -s "http://192.168.1.12:9200/_cat/indices?v"
health status index          uuid                   pri rep docs.count docs.deleted store.size pri.store.size
dataset.size
yellow open   neurapent_loot iE3RsWwVRr-6n8pP_7SZ2g   1   1          5            0       13kb           13kb
13kb
```

BUSINESS IMPACT

An attacker can connect to the Elasticsearch instance and read all data from all indices, including the suggestively named 'neurapent_loot' index. They can also modify or delete all data, leading to a complete data breach and/or data loss. Furthermore, an attacker could potentially write large amounts of data to cause a denial of service by exhausting disk space.

REMEDIATION

Enable and enforce authentication on the Elasticsearch cluster. In the `elasticsearch.yml` configuration file, ensure security features are enabled by setting `xpack.security.enabled: true`. Additionally, configure transport layer security (TLS) for both HTTP and transport layers by setting `xpack.security.http.ssl.enabled: true` and `xpack.security.transport.ssl.enabled: true`. Access to the Elasticsearch API should also be restricted at the network level by a firewall to only allow connections from trusted application servers.

F-004 Attack Chain: Unauthenticated Redis to RCE and Internal Pivot

SEVERITY	CVSS V3.1	STATUS	ASSET
CRITICAL	9.8	Potential	192.168.1.12:6379/redis

DESCRIPTION

An attacker can connect to the exposed Redis instance without authentication. By leveraging standard Redis commands, the attacker can write files to arbitrary locations on the filesystem, leading to Remote Code Execution by planting a web shell in the Apache document root. This initial foothold can then be used to pivot and attack other internal services like MariaDB, which currently blocks external connections.

AFFECTED ASSET

Host	192.168.1.12
Port	6379
Service	redis (Redis key-value store 7.4.8)

EVIDENCE

```
Chain: Unauthenticated Redis Access -> Apache Server Configuration Disclosure -> Implied Action -> Implied Action -> Empty/No Password
```

BUSINESS IMPACT

An attacker can connect to the exposed Redis instance without authentication. By leveraging standard Redis commands, the attacker can write files to arbitrary locations on the filesystem, leading to Remote Code Execution by planting a web shell in the Apache document root. This initial foothold can then be used to pivot and attack other internal services like MariaDB, which currently blocks external connections.

REMEDIATION

Update port 6379 to latest version; Review vendor security advisories; Implement network segmentation; Monitor for exploitation attempts

F-005 Writable Network Share

SEVERITY	STATUS	ASSET
HIGH	Potential	192.168.1.12:139/?

DESCRIPTION

A network share allows write access, potentially enabling malware deployment or data tampering.

AFFECTED ASSET

Host	192.168.1.12
Port	139

Service	— (Microsoft Windows netbios-ssn)
Detection Tool	nuclei

EVIDENCE


```
{"template":"javascript/cves/2025/CVE-2025-46817.yaml","template-
url":"https://cloud.projectdiscovery.io/public/CVE-2025-46817","template-id":"CVE-2025-46817","template-
path":"/root/nuclei-templates/javascript/cves/2025/CVE-2025-46817.yaml","info":{"name":"Redis \u003c 8.2.1
lua script - Integer Overflow","author":["pussycat0x"],"tags":
["cve","cve2025","js","redis","network","passive","authenticated","vuln"],"description":"Redis is an open
source, in-memory database that persists on disk. Versions 8.2.1 and below allow an authenticated user to use
a specially crafted Lua script to cause an integer overflow and potentially lead to remote code execution The
problem exists in all versions of Redis with Lua scripting. This issue is fixed in version
8.2.2.\n","impact":"Authenticated users can execute arbitrary code remotely, potentially compromising the
entire system.\n","reference":["https://github.com/dwisiswant0/CVE-2025-
46817/blob/master/README.md"],"severity":"critical","metadata":{"shodan-
query":"product:\"redis\"","verified":true,"max-request":1},"remediation":"Update to version 8.2.2 or
later.\n"},"extractor-
name":"version","type":"javascript","host":"192.168.1.12","port":"6379","url":"192.168.1.12:6379","matched-
at":"192.168.1.12:6379","extracted-results":["7.4.8"],"request":"const redis =
require('nuclei/redis');\nconst info = redis.GetServerInfo(Host, Port);\nExport(info);","response":"#
Server\r\nredis_version:7.4.8\r\nredis_git_sha1:00000000\r\nredis_git_dirty:0\r\nredis_build_id:61036196736d9fd9\
6.6.114.1-microsoft-standard-WSL2 x86_64\r\narch_bits:64\r\nmonotonic_clock:POSIX
```


BUSINESS IMPACT

Data integrity risk. Attackers can upload malicious files or modify existing data.

REMEDIATION

Review and restrict share permissionsImplement least-privilege accessEnable auditing on sensitive sharesUse network segmentation

F-006 CVE-2023-38709

SEVERITY	CVSS V3.1	STATUS	CVE	ASSET
HIGH	7.3	Potential	CVE-2023-38709	192.168.1.12:80/http

DESCRIPTION

Potential vulnerability detected: CVE-2023-38709

AFFECTED ASSET

Host	192.168.1.12
Port	80
Service	http (Apache httpd 2.4.58 ((Win64) OpenSSL/3.1.3 PHP/8.2.12))
CVE Reference	CVE-2023-38709
Detection Tool	nmap

EVIDENCE

```
$ nmap -sV --script vulners -p 80 192.168.1.12
|       CVE-2025-49812  7.4 https://vulners.com/cve/CVE-2025-49812
|       CVE-2026-29168  7.3 https://vulners.com/cve/CVE-2026-29168
|       CVE-2023-38709  7.3 https://vulners.com/cve/CVE-2023-38709
|       CNVD-2024-36395 7.3 https://vulners.com/cnvd/CNVD-2024-36395
|       CVE-2026-33523  6.5 https://vulners.com/cve/CVE-2026-33523
```

BUSINESS IMPACT

Potential vulnerability detected: CVE-2023-38709

REMEDIATION

Update web server to latest version; Apply all security patches; Implement WAF protection; Review application code for vulnerabilities

F-007 CVE-2025-49812

SEVERITY	CVSS V3.1	STATUS	CVE	ASSET
HIGH	7.4	Potential	CVE-2025-49812	192.168.1.12:443/ssl/http

DESCRIPTION

Potential vulnerability detected: CVE-2025-49812

AFFECTED ASSET

Host	192.168.1.12
Port	443
Service	ssl/http (Apache httpd 2.4.58 ((Win64) OpenSSL/3.1.3 PHP/8.2.12))
CVE Reference	CVE-2025-49812
Detection Tool	nmap

EVIDENCE

```
$ nmap -sV --script vulners -p 443 192.168.1.12
|       45D138AD-BEC6-552A-91EA-8816914CA7F4    7.5 https://vulners.com/githubexploit/45D138AD-BEC6-552A-91EA-8816914CA7F4    *EXPLOIT*
|       0E08753E-C6D7-5E76-A61F-6CA6F7F87AA8    7.5 https://vulners.com/githubexploit/0E08753E-C6D7-5E76-A61F-6CA6F7F87AA8    *EXPLOIT*
|       CVE-2025-49812  7.4 https://vulners.com/cve/CVE-2025-49812
|       CVE-2026-29168  7.3 https://vulners.com/cve/CVE-2026-29168
|       CVE-2023-38709  7.3 https://vulners.com/cve/CVE-2023-38709
```

BUSINESS IMPACT

Potential vulnerability detected: CVE-2025-49812

REMEDIATION

Update port 443 to latest version; Review vendor security advisories; Implement network segmentation; Monitor for exploitation attempts

F-008 CORS Misconfiguration Allows Credentialed Cross-Origin Data Theft

SEVERITY	CVSS V3.1	STATUS	ASSET
HIGH	6.9	Potential	192.168.1.12:5000/http

DESCRIPTION

The application's Cross-Origin Resource Sharing (CORS) policy improperly validates the 'Origin' header. The server reflects the 'null' origin in the 'Access-Control-Allow-Origin' header and simultaneously sets 'Access-Control-Allow-Credentials' to 'true'. A 'null' origin can be generated by sandboxed documents, local HTML files, or certain redirects. This combination allows a malicious website to make authenticated requests to the API on behalf of a user and read the sensitive response.

AFFECTED ASSET

Host	192.168.1.12
Port	5000
Service	http (Uvicorn)

EVIDENCE

```
$ nuclei -u http://192.168.1.12:5000 -t http/vulnerabilities/generic/cors-misconfig.yaml
{
  "request": "GET / HTTP/1.1\r\nHost: 192.168.1.12:5000\r\nOrigin: null\r\n",
  "response": "HTTP/1.1 200 OK\r\n...\r\nAccess-Control-Allow-Credentials: true\r\nAccess-Control-Allow-
Origin: null\r\n...\r\n",
  "matcher-name": "arbitrary-origin"
}
```

BUSINESS IMPACT

An attacker can craft a malicious webpage. When a user who is authenticated to the 'Enterprise LLM Pentesting Gateway' visits this page, the attacker's script can make API calls as the user to endpoints like '/v1/chat/completions' or '/v1/autopilot/run'. The script can then exfiltrate the sensitive data returned by the API, potentially leading to data theft, unauthorized actions, and full account compromise.

REMEDIATION

The CORS policy should be configured with a strict whitelist of trusted origins. Do not reflect arbitrary origins like 'null' in the 'Access-Control-Allow-Origin' header, especially when 'Access-Control-Allow-Credentials' is 'true'. If 'null' origin support is absolutely required, ensure 'Access-Control-Allow-Credentials' is set to 'false' for those requests. A more secure approach is to disallow 'null' origins in production environments entirely.

F-009 Intel Active Management Technology (AMT) Vulnerable to Privilege Escalation

SEVERITY	CVSS V3.1	STATUS	CVE	ASSET
HIGH	8.2	Potential	CVE-2023-28465	192.168.1.12:623/unknown

DESCRIPTION

The version of Intel Active Management Technology (AMT) detected (12.0.93.2331) is outdated and vulnerable to multiple security issues, as detailed in Intel Security Advisory INTEL-SA-00923. These vulnerabilities (including CVE-2023-28465, CVE-2023-28466, CVE-2023-28467, and CVE-2023-28468) can allow an unauthenticated user with network access to escalate privileges on the affected system.

AFFECTED ASSET

Host	192.168.1.12
Port	623
Service	unknown (Intel Active Management Technology User Notification Service httpd 12.0.93.2331)
CVE Reference	CVE-2023-28465

EVIDENCE

```
$ nmap -sV -p 623 192.168.1.12
PORT    STATE SERVICE VERSION
623/tcp open  http    Intel Active Management Technology User Notification Service httpd 12.0.93.2331
|_http-server-header: Intel(R) Active Management Technology 12.0.93.2331
Service Info: CPE: cpe:/h:intel:active_management_technology:12.0.93.2331
```

BUSINESS IMPACT

An attacker on the same network segment could exploit these vulnerabilities to gain higher privileges on the system. Since AMT operates at a hardware level, a successful exploit could lead to a complete compromise of the host, bypassing operating system-level security controls and potentially gaining persistent access.

REMEDIATION

The system's firmware/BIOS should be updated to a version that includes Intel AMT firmware version 12.0.94 or later. This update must be obtained from the system or motherboard manufacturer (e.g., Dell, HP, Lenovo). If AMT functionality is not required for remote management, it is strongly recommended to disable it entirely in the system's BIOS/UEFI settings to reduce the attack surface.

F-010 Writable Network Share

SEVERITY	CVSS V3.1	STATUS	ASSET
HIGH	9.8	Potential	192.168.1.12:139/multi-service

DESCRIPTION

A network share allows write access, potentially enabling malware deployment or data tampering.

AFFECTED ASSET

Host	192.168.1.12
Port	139
Service	multi-service (Microsoft Windows netbios-ssn)

EVIDENCE

```
Chain: Writable Network Share -> Apache Server Configuration Disclosure -> Implied Action -> Implied Action
```

BUSINESS IMPACT

Data integrity risk. Attackers can upload malicious files or modify existing data.

REMEDIATION

Review and restrict share permissionsImplement least-privilege accessEnable auditing on sensitive sharesUse network segmentation

F-011 Weak TLS Cipher Suites Supported (Logjam)

SEVERITY	CVSS V3.1	STATUS	CVE	ASSET
MEDIUM	3.0	Potential	CVE-2015-4000	192.168.1.12:443/http

DESCRIPTION

The server supports TLS cipher suites that use Diffie-Hellman (DH) key exchange with weak 1024-bit parameters. This makes the server vulnerable to the Logjam attack (CVE-2015-4000), where a man-in-the-middle attacker can downgrade the connection to export-grade cryptography and decrypt the TLS traffic.

AFFECTED ASSET

Host	192.168.1.12
Port	443
Service	http (Apache httpd 2.4.58 ((Win64) OpenSSL/3.1.3 PHP/8.2.12))

CVE Reference	CVE-2015-4000

EVIDENCE

```
$ nmap -p 443 --script ssl-enum-ciphers 192.168.1.12
| ssl-enum-ciphers:
|   TLSv1.2:
|     ciphers:
|       TLS_DHE_RSA_WITH_AES_256_GCM_SHA384 (dh 1024) - F
|       TLS_DHE_RSA_WITH_AES_128_GCM_SHA256 (dh 1024) - F
|_  least strength: F
```

BUSINESS IMPACT

An attacker with the ability to intercept network traffic (man-in-the-middle) can downgrade the secure connection and decrypt sensitive information exchanged between the user and the server, such as login credentials, session cookies, and personal data.

REMEDIATION

Disable all DHE cipher suites and ensure a strong DH parameter file (at least 2048 bits) is used. In Apache's `ssl.conf` or equivalent, modify the `SSLCipherSuite` directive to exclude DHE ciphers and prioritize modern, secure ciphers. Generate a new DH parameter file: openssl dhparam -out /etc/ssl/certs/dhparam.pem 2048 Then, add this line to your Apache SSL configuration: SSLOpenSSLConfCmd DHParameters "/etc/ssl/certs/dhparam.pem"

F-012 Sensitive API Endpoints and Service Details Disclosed on Root Endpoint

SEVERITY	CVSS V3.1	STATUS	ASSET
MEDIUM	6.5	Potential	192.168.1.12:5000/http

DESCRIPTION

The root endpoint ('/') of the web server returns a JSON object containing a full list of available API endpoints, the service name ("Enterprise LLM Pentesting Gateway"), and its version. This information provides a complete roadmap of the application's functionality to any unauthenticated user, eliminating the need for reconnaissance.

AFFECTED ASSET

Host	192.168.1.12
Port	5000
Service	http (Uvicorn)

EVIDENCE

```
$ curl http://192.168.1.12:5000/
{
  "service": "Enterprise LLM Pentesting Gateway",
  "version": "3.0",
  "endpoints": {
    "/ping": "Status check",
    "/health": "Health check",
    "/v1/chat/completions": "AI Gateway (POST)",
    "/v1/autopilot/run": "Autopilot (POST)",
    "/run/{tool_name}": "Tool execution (POST)"
  }
}
```

BUSINESS IMPACT

An attacker can immediately understand the purpose of the application and identify all available endpoints to probe for further vulnerabilities. This accelerates targeted attacks against sensitive functions like the AI gateway

('/v1/chat/completions') and remote execution features ('/v1/autopilot/run', '/run/{tool_name}').

REMEDIATION

The root endpoint should not publicly disclose the application's entire API structure. This index page should be removed or protected by authentication and authorization controls, making it available only to privileged users or developers. Return a generic 404 Not Found or an empty page for the root path.

F-013 Intel Active Management Technology (AMT) Version Disclosure

SEVERITY	CVSS V3.1	STATUS	ASSET
MEDIUM	4.2	Potential	192.168.1.12:16992/http

DESCRIPTION

The web server running on port 16992 identifies itself as "Intel(R) Active Management Technology 12.0.93.2331". Intel AMT is a powerful out-of-band management technology that can provide hardware-level access to a system. Disclosing the exact version number allows an attacker to search for specific vulnerabilities and exploits, significantly reducing the effort required to find a viable attack vector.

AFFECTED ASSET

Host	192.168.1.12
Port	16992
Service	http (Intel Active Management Technology User Notification Service httpd 12.0.93.2331)

EVIDENCE

```
$ nmap -sV -p 16992 192.168.1.12
PORT      STATE SERVICE VERSION
16992/tcp open  http    Intel Active Management Technology User Notification Service httpd 12.0.93.2331
|_http-server-header: Intel(R) Active Management Technology 12.0.93.2331
Service Info: CPE: cpe:/h:intel:active_management_technology:12.0.93.2331
```

BUSINESS IMPACT

An attacker can use the disclosed version number (12.0.93.2331) to identify and launch targeted attacks against known vulnerabilities in this specific version of Intel AMT. A successful exploit could grant the attacker complete control over the system at the hardware level, bypassing all operating system security controls. This could lead to remote code execution, data theft, or persistent compromise of the host.

REMEDIATION

The Intel AMT firmware should be updated to the latest version available from the system manufacturer to mitigate known vulnerabilities. Additionally, access to the AMT management ports (e.g., 16992, 16993) should be restricted at the network level, allowing connections only from trusted management subnets or administrative hosts. If AMT is not required for remote administration, it should be fully disabled in the system's BIOS/UEFI settings.

F-014 VMware Authentication Daemon Version Disclosure Indicates End-of-Life Software

SEVERITY	CVSS V3.1	STATUS	ASSET
MEDIUM	3.0	Potential	192.168.1.12:902/ssl/vmware-auth

DESCRIPTION

The service running on port 902 identified itself as 'VMware Authentication Daemon 1.10'. This version is associated with very old VMware products (e.g., VMware Server 1.x, Workstation 6.x, ESX 3.x) which are long past their end-of-life and no longer receive security updates. Running unsupported, EOL software is a significant security risk as it likely contains unpatched vulnerabilities.

AFFECTED ASSET

Host	192.168.1.12
Port	902
Service	ssl/vmware-auth (VMware Authentication Daemon 1.10 (Uses VNC, SOAP))

EVIDENCE

```
$ nmap -sV -p 902 192.168.1.12
PORT    STATE SERVICE         VERSION
902/tcp open  ssl/vmware-auth VMware Authentication Daemon 1.10 (Uses VNC, SOAP)
```

BUSINESS IMPACT

An attacker can use this precise version number to find and use publicly available exploits. Given the age of the software, vulnerabilities could range from denial of service to authentication bypass, granting unauthorized access to the VMware management plane and the virtual machines it controls.

REMEDIATION

The underlying VMware product is end-of-life and must be upgraded to a modern, supported version (e.g., a recent release of VMware ESXi). If this service is part of a larger, legacy application, consider migrating the application to a supported platform. If the service cannot be upgraded, access to this port should be strictly restricted by a firewall to only allow connections from trusted management IP addresses.

F-015 VMware Authentication Daemon Version 1.0 Information Disclosure

SEVERITY	CVSS V3.1	STATUS	ASSET
MEDIUM	4.2	Potential	192.168.1.12:912/vmware-auth

DESCRIPTION

The VMware Authentication Daemon (vmauthd) service running on port 912 is broadcasting a banner that reveals its name, exact version (1.0), and the protocols it uses (SOAP, VNC). Version 1.0 is extremely old, dating back to legacy VMware products from the mid-2000s (e.g., VMware Workstation 5.x/6.x, VMware Server 1.x). Such legacy software is unsupported and almost certainly contains numerous publicly known vulnerabilities.

AFFECTED ASSET

Host	192.168.1.12
Port	912
Service	vmware-auth (VMware Authentication Daemon 1.0 (Uses VNC, SOAP))

EVIDENCE

```
$ nc -nv -w 1 192.168.1.12 912
220 VMware Authentication Daemon Version 1.0, ServerDaemonProtocol:SOAP, MKSDisplayProtocol:VNC , , ,

(UNKNOWN) [192.168.1.12] 912 (?) open
```

BUSINESS IMPACT

An attacker can use this precise version information to find and launch targeted exploits for known vulnerabilities in VMware Authentication Daemon 1.0. Publicly known vulnerabilities for this service could lead to denial of service, information disclosure (e.g., arbitrary file reads via directory traversal), or potentially remote code execution.

REMEDIATION

The VMware Authentication Daemon is a component of a larger VMware product (e.g., VMware Workstation, Server, ESX/ESXi). The underlying product should be identified and upgraded to a modern, supported version. If the service is not essential for business operations, it should be disabled or firewalled to restrict access to only trusted hosts. On a Linux host, the service can often be stopped and disabled using the system's service manager (e.g., `sudo systemctl stop vmware.service && sudo systemctl disable vmware.service` or `/etc/init.d/vmware stop`).

F-016 Outdated MariaDB Version Detected

SEVERITY	CVSS V3.1	STATUS	ASSET
MEDIUM	3.0	Potential	192.168.1.12:3306/mysql

DESCRIPTION

The service is running MariaDB version 10.3.23 or an earlier release. MariaDB 10.3 reached its End-of-Life (EOL) in May 2023. Software that is past its EOL no longer receives security patches or bug fixes from the vendor, making it a potential target for attackers leveraging newly discovered vulnerabilities.

AFFECTED ASSET

Host	192.168.1.12
Port	3306
Service	mysql (MariaDB 10.3.23 or earlier (unauthorized))

EVIDENCE

```
$ nmap -sV -p 3306 192.168.1.12
PORT     STATE SERVICE VERSION
3306/tcp open  mysql   MariaDB 10.3.23 or earlier (unauthorized)
```

BUSINESS IMPACT

An attacker can use the specific version information to search for and exploit publicly known vulnerabilities affecting MariaDB 10.3 and older. Depending on the available exploits, this could lead to denial of service, information disclosure, or remote code execution. The lack of vendor support means any future vulnerabilities discovered will remain unpatched.

REMEDIATION

Upgrade the MariaDB instance to a currently supported stable version (e.g., 10.6, 10.11, or newer). Before performing the upgrade, ensure a complete backup of all databases is taken. Follow the official MariaDB upgrade documentation for the specific version path. 1. Perform a full backup of all databases: `mysqldump --all-databases > /path/to/backup.sql` 2. Stop the MariaDB service. 3. Follow the vendor's instructions to uninstall the old version and install the new version. 4. Run `mysql_upgrade` after installation to check and update system tables. 5. Restart the service and verify application functionality.

F-017 AnyDesk Remote Desktop Service Identified

SEVERITY	CVSS V3.1	STATUS	ASSET
MEDIUM	3.3	Potential	192.168.1.12:7070/ssl/realserver?

DESCRIPTION

The SSL certificate presented by the service on port 7070 identifies it as "AnyDesk Client". AnyDesk is a remote desktop

application that allows for remote control of the host machine. The certificate is self-signed and has an unusually long validity period, which is characteristic of a default AnyDesk installation. The presence of remote access software significantly increases the host's attack surface.

AFFECTED ASSET

Host	192.168.1.12
Port	7070
Service	ssl/realserver? (| ssl-cert: Subject: commonName=AnyDesk Client)

EVIDENCE

```
$ nmap -sV --script ssl-cert -p 7070 192.168.1.12
| ssl-cert: Subject: commonName=AnyDesk Client
| Issuer: commonName=AnyDesk Client
| Not valid before: 2025-12-13T09:37:23
|_Not valid after:  2075-12-01T09:37:23
```

BUSINESS IMPACT

An attacker can focus their efforts on compromising the AnyDesk service. If successful, for example by guessing a weak password for unattended access or exploiting a vulnerability, they could gain complete remote control of the host operating system.

REMEDIATION

Review the business need for AnyDesk on this host. If it is not required, uninstall the software. If it is required, ensure it is configured securely: 1. Disable unattended access if not strictly necessary. 2. If unattended access is enabled, protect it with a strong, unique password. 3. Use an access control list (ACL) in the AnyDesk settings to restrict which AnyDesk IDs can connect. 4. Keep the AnyDesk client updated to the latest version to protect against known vulnerabilities.

F-018 Neo4j Service Leaks Version and Internal IP Address

SEVERITY	CVSS V3.1	STATUS	ASSET
MEDIUM	6.5	Potential	192.168.1.12:7474/unknown

DESCRIPTION

The Neo4j service running on port 7474 responds to unauthenticated HTTP GET requests with a JSON object containing sensitive information. This includes the exact software version (5.26.23), the edition (community), and an internal IP address (172.18.0.4) which appears to belong to a container or internal network.

AFFECTED ASSET

Host	192.168.1.12
Port	7474
Service	unknown (| fingerprint-strings:)

EVIDENCE

```
$ nmap -p 7474 --script=fingerprint-strings 192.168.1.12
|   GetRequest:
|     HTTP/1.1 200 OK
|     Content-Type: application/json
|
{"bolt_routing":"neo4j://172.18.0.4:7687","query":"http://172.18.0.4:7474/db/{databaseName}/query/v2","transactio
```

BUSINESS IMPACT

An attacker can use the precise version number to find known vulnerabilities and public exploits for this specific version of Neo4j. The leaked internal IP address reveals details about the internal network architecture (e.g., the use of containers), which aids an attacker in mapping the internal network and planning lateral movement.

REMEDIATION

Configure Neo4j to not expose sensitive information to unauthenticated users. In the `neo4j.conf` file, ensure that the advertised address is set correctly to a public-facing address, not an internal one, using the `server.discovery.advertised_address` setting. Additionally, implement access controls or a web application firewall (WAF) to restrict access to the root API endpoint or filter out sensitive headers and body content.

F-019 Neo4j Service Configured with Insecure CORS Policy

SEVERITY	CVSS V3.1	STATUS	ASSET
MEDIUM	3.8	Potential	192.168.1.12:7474/unknown

DESCRIPTION

The Neo4j HTTP API is configured with a permissive Cross-Origin Resource Sharing (CORS) policy (`Access-Control-Allow-Origin: *`). This allows web pages from any domain to make requests to the Neo4j API, bypassing the browser's Same-Origin Policy.

AFFECTED ASSET

Host	192.168.1.12
Port	7474
Service	unknown (| fingerprint-strings:)

EVIDENCE

```
$ nmap -p 7474 --script=fingerprint-strings 192.168.1.12
|   GetRequest:
|     HTTP/1.1 200 OK
|     Access-Control-Allow-Origin: *
```

BUSINESS IMPACT

An attacker can create a malicious website that, when visited by a user, makes requests to the Neo4j database from the user's browser. If the database allows anonymous access or if the user is authenticated to the service, the attacker could potentially read, modify, or exfiltrate data from the database on behalf of the victim user.

REMEDIATION

Restrict the CORS policy to only allow trusted domains. In the `neo4j.conf` file, set `dbms.security.http_access_control_allow_origin` to a specific whitelist of domains (e.g., `dbms.security.http_access_control_allow_origin=https://app.yourdomain.com`). If no cross-domain access is required from a web browser, this feature should be disabled.

F-020 Insecure CORS Policy Allows Cross-Domain Data Access

SEVERITY	CVSS V3.1	STATUS	ASSET
MEDIUM	3.8	Potential	192.168.1.12:7687/bolt?

DESCRIPTION

The service on this port, identified as an HTTP-based API, is configured with a permissive Cross-Origin Resource Sharing (CORS) policy. The 'Access-Control-Allow-Origin: *' header instructs web browsers to allow scripts from any website to

make requests to this service and read the responses. This completely bypasses the browser's Same-Origin Policy, which is a critical security control.

AFFECTED ASSET

Host	192.168.1.12
Port	7687
Service	bolt? (1 service unrecognized despite returning data. If you know the service/version, please submit the following fingerprint at https://nmap.org/cgi-bin/submit.cgi?new-service :)

EVIDENCE

```
$ nmap -sV -p 7687 192.168.1.12
| fingerprint-strings:
|   FourOhFourRequest, GetRequest:
|     HTTP/1.0 200 OK
|     content-type: application/json
|     access-control-allow-origin: *
```

BUSINESS IMPACT

An attacker can create a malicious website. If a user on the same network as the target visits this website, the attacker's script can make requests to this API (http://192.168.1.12:7687) from the user's browser. The script can then read any data returned by the API, potentially exfiltrating sensitive information. If the API allows state-changing operations (e.g., POST, DELETE), the attacker could also perform actions on behalf of the user.

REMEDIATION

The CORS policy should be restricted to a whitelist of trusted domains. If the API is not intended to be accessed by any browser-based clients, remove the `Access-Control-Allow-Origin` header entirely. If access is required from specific web applications, configure the server to return the specific origin, for example: `Access-Control-Allow-Origin: https://trusted.app.com`.

F-021 Elasticsearch Version and Cluster Information Disclosure

SEVERITY	CVSS V3.1	STATUS	ASSET
MEDIUM	3.3	Potential	192.168.1.12:9200/elasticsearch

DESCRIPTION

The Elasticsearch service banner discloses its exact version (8.12.2), node name (68930e1807f4), cluster name (docker-cluster), and underlying Lucene version (9.9.2). This information helps an attacker identify potential vulnerabilities and tailor their attacks.

AFFECTED ASSET

Host	192.168.1.12
Port	9200
Service	elasticsearch (Elasticsearch REST API 8.12.2 (name: 68930e1807f4; cluster: docker-cluster; Lucene 9.9.2))

EVIDENCE

```
$ nmap -sV -p 9200 192.168.1.12
PORT     STATE SERVICE VERSION
9200/tcp open  http    Elasticsearch REST API 8.12.2 (name: 68930e1807f4; cluster: docker-cluster; Lucene
9.9.2)
```

BUSINESS IMPACT

An attacker can use the disclosed version information to search for known public exploits targeting Elasticsearch 8.12.2. This significantly reduces the effort required to find a viable attack vector and provides context about the environment (e.g., 'docker-cluster').

REMEDIATION

The primary remediation is to prevent unauthenticated access as detailed in the 'Unauthenticated Elasticsearch API Access' finding. Once authentication is in place, this information is only available to authenticated users. Additionally, restrict network access to the Elasticsearch port using a firewall, allowing connections only from trusted sources.

F-022 Verbose Server and Technology Banners Disclosed

SEVERITY		STATUS	ASSET
LOW	0	Potential	192.168.1.12:80/http

DESCRIPTION

The HTTP response headers reveal detailed version information for the web server stack, including Apache, OpenSSL, and PHP. This information is also present on default error pages.

AFFECTED ASSET

Host	192.168.1.12
Port	80
Service	http (Apache httpd 2.4.58 ((Win64) OpenSSL/3.1.3 PHP/8.2.12))

EVIDENCE

```
$ curl -I http://192.168.1.12
Server: Apache/2.4.58 (Win64) OpenSSL/3.1.3 PHP/8.2.12
X-Powered-By: PHP/8.2.12
```

BUSINESS IMPACT

An attacker can use the precise version numbers to quickly find and tailor exploits for any known vulnerabilities in these components, reducing the effort required for a successful attack.

REMEDIATION

Configure the web server to suppress detailed version information. In the Apache `httpd.conf` file, set `ServerTokens Prod` and `ServerSignature Off`. In the `php.ini` file, set `expose_php = Off`. Restart the Apache service after making these changes.

F-023 Missing Subresource Integrity (SRI)

SEVERITY		STATUS	ASSET
LOW	0	Potential	192.168.1.12:80/http

DESCRIPTION

The application loads external JavaScript and CSS files from a third-party CDN without using the Subresource Integrity (SRI) attribute. This allows the CDN to deliver malicious code to users if it is ever compromised.

AFFECTED ASSET

Host	192.168.1.12

Port	80
Service	http (Apache httpd 2.4.58 ((Win64) OpenSSL/3.1.3 PHP/8.2.12))

EVIDENCE

```
$ nuclei -t http/misconfiguration/missing-sri.yaml -u http://192.168.1.12
[missing-sri] [http] [info] http://192.168.1.12:80/ [extracted-results: //code.jquery.com/jquery-
1.10.2.min.js, //cdnjs.cloudflare.com/ajax/libs/font-awesome/3.1.0/css/font-awesome.min.css]
```

BUSINESS IMPACT

If the third-party CDN is compromised, an attacker could inject malicious code into the included files. This code would then execute in the browsers of the application's users, potentially leading to credential theft, session hijacking, or other client-side attacks.

REMEDIATION

Add the `integrity` and `crossorigin` attributes to all `<script>` and `<link>` tags that load resources from external domains. The hash can be generated using a tool like `openssl` or an online generator. First, upgrade the outdated libraries, then apply SRI. Example for an upgraded jQuery: `<script src="https://code.jquery.com/jquery-3.7.1.min.js" integrity="sha256-eKhayi8LEQwp4NKxN+CfCh+3qOVUtJn3QNZ0TciWLP4=" crossorigin="anonymous"></script>`

F-024 Verbose HTTP Server Header Discloses Technology

SEVERITY		STATUS	ASSET
LOW	0	Potential	192.168.1.12:5000/http

DESCRIPTION

The 'Server' HTTP header returned by the application reveals that the service is running 'uvicorn'. This information helps an attacker build a profile of the technologies used by the target.

AFFECTED ASSET

Host	192.168.1.12
Port	5000
Service	http (Uvicorn)

EVIDENCE

```
$ nmap -sV -p 5000 192.168.1.12
PORT     STATE SERVICE VERSION
5000/tcp open  http    Uvicorn
|_http-server-header: uvicorn
```

BUSINESS IMPACT

An attacker can use the disclosed server technology ('uvicorn') to search for specific vulnerabilities affecting that software. If a new vulnerability is discovered in Uvicorn, the attacker knows this system is a potential target, reducing the effort needed to find an exploit.

REMEDIATION

Configure the web server to suppress or obfuscate the 'Server' header. For Uvicorn, this is often handled by a reverse proxy like Nginx or by creating a middleware in the application (e.g., in FastAPI) to remove or modify the header before the response is sent to the client.

7. Attack Narratives

The following multi-step attack chains were identified and verified during the engagement. Each narrative describes a realistic exploitation path from initial access to material business impact, in the order an attacker would have executed them.

NARRATIVE 01 — UNAUTHENTICATED REDIS TO RCE AND INTERNAL PIVOT

SEVERITY	STATUS
CRITICAL	IN_PROGRESS

FINAL IMPACT	STEPS
Full system compromise via Remote Code Execution, and subsequent access to internal database services.	5

An attacker can connect to the exposed Redis instance without authentication. By leveraging standard Redis commands, the attacker can write files to arbitrary locations on the filesystem, leading to Remote Code Execution by planting a web shell in the Apache document root. This initial foothold can then be used to pivot and attack other internal services like MariaDB, which currently blocks external connections.

REPRODUCTION STEPS

Unauthenticated Redis Access — Connect to the Redis service on port 6379 without authentication to gain command execution capabilities within the database.

Apache Server Configuration Disclosure — Query the Apache server on port 80 to retrieve the `/server-info` page, identifying the web server's document root directory.

Implied Action — Use Redis commands (e.g., CONFIG SET dir, CONFIG SET dbfilename, SAVE) to write a PHP web shell into the discovered Apache web root.

Implied Action — Access the uploaded web shell via a browser to achieve Remote Code Execution on the target host.

Empty/No Password — From the compromised host, connect to the MariaDB instance on port 3306 from localhost, bypassing the network ACL that previously blocked access.

NARRATIVE 02 — WRITABLE SMB SHARE TO WEB SHELL RCE

SEVERITY	STATUS	FINAL IMPACT	STEPS
CRITICAL	IN_PROGRESS	Full system compromise via Remote Code Execution.	4

An attacker can abuse a misconfigured writable SMB share to place a web shell into the document root of the Apache web server. The path for the web root can be confirmed via a separate information disclosure vulnerability. Executing this web shell provides the attacker with Remote Code Execution on the system.

REPRODUCTION STEPS

Writable Network Share — Connect to the writable network share on port 139 with guest or null credentials.

Apache Server Configuration Disclosure — Use the information disclosure on port 80 to confirm the exact file path of the web root, verifying it is accessible via the SMB share.

Implied Action — Upload a PHP web shell payload to the web root directory via the writable SMB share.

Implied Action — Execute the web shell by requesting the corresponding URL from the Apache server on port 80, gaining RCE.

8. Conclusion & Recommendations

The engagement against **192.168.1.12** identified 24 security findings, including **4 Critical** and **6 High**. The overall security posture of the in-scope environment is rated **Critical**.

The customer is advised to remediate findings in the following priority order:

Address all **Critical** findings within 7 days. These represent immediate, exploitable risk.

Address all **High** findings within 30 days.

Address Medium findings within the next maintenance window (typically 90 days).

Treat Low and Informational findings as hardening opportunities.

NeuraPent recommends a re-test be conducted after remediation to verify that controls are effective and no regression has been introduced. Periodic continuous testing — recommended quarterly for production environments — will help maintain a strong security posture as the threat landscape evolves.

— END OF REPORT —

NPR-20260601-E5A7FF · 2026-06-01 · NEURAPENT AUTONOMOUS SECURITY