

Chapter 3: Vulnerability Assessment

A vulnerability identified by an automated tool serves as an initial indicator rather than a definitive conclusion. Its significance is established only when an analyst contextualizes it with device-specific evidence and determines the appropriate response.

Learning Objectives

By the end of this chapter, students should be able to:

1. Explain vulnerability assessment as detective analytics rather than scanner output collection.
2. Distinguish a candidate vulnerability, a confirmed weakness, a false positive, and a prioritized remediation item.
3. Explain how CPE precision controls NVD results.
4. Describe Breakwater's lookup path: exact cpeName, virtualMatchString fallback, Redis caching, deduplication, backoff, and per-host summaries.
5. Interpret Common Vulnerability Scoring System (CVSS) version 3.1 base vectors as structured evidence, not just numeric scores.
6. Compare NVD lookups, OpenVAS, Nuclei, and default-credential testing by method, confidence, intrusiveness, and failure mode.
7. Explain the default credential checker's protocol mapping, vendor-first credential selection, and false-positive guardrails.
8. Use validation records to suppress known false-positive CVE/device pairs without deleting the original evidence history.
9. Correlate findings across tools without double-counting.
10. Build a remediation queue that combines severity, evidence strength, exposure, device role, and consequence.

Problem Framing

The primary challenge in assessment extends beyond tool execution; it involves mastering evidence interpretation, drawing well-supported conclusions, and recognizing the limits of current knowledge.

The code can query the National Vulnerability Database, run OpenVAS, invoke Nuclei, test public factory credentials, and record validation decisions. These outputs are valuable only if the analyst maintains them as separate evidence classes. Passive product-name matches, active template hits, Greenbone results, successful logins, and firmware-to-CVE rules represent different evidence. They should not be combined into a single count before their distinctions are clear.

For Chapter 3, the evidence model has five parts:

Observation	CPEs, CVE records, CVSS vectors, OpenVAS results, Nuclei findings, credential outcomes
-------------	--

Chapter 3: Vulnerability Assessment

State	Per-host evidence with source, method, confidence, and validation status
Decision	Candidate, confirmed, false positive, risk accepted, or unresolved
Guardrail	Do not escalate from candidate to fact without applicable evidence
Action handoff	Prioritized queue item, validation task, or documented uncertainty

Table 3.1. Evidence model for vulnerability assessment.

The current codebase implements some of this process directly and leaves some to policy. It stores validation records, can suppress false-positive CVE/device pairs, and keeps active-tool findings in separate host fields. It emits phase progress events to indicate scan status. The system does not automatically prioritize high CVSS scores over lower ones; such decisions depend on the evidence model and operational context.

Caution: The techniques described in this chapter have operational boundaries. Always verify assumptions against the specific deployment environment before relying on any automated output.

What did the system observe, and what conclusion is it allowed to draw from that observation?

This approach mitigates the risks associated with automation. The effectiveness of a system is determined by its ability to preserve the context and rationale behind each piece of evidence, rather than the length of its vulnerability list. This ensures that future analysts can trace the reasoning underlying every conclusion.

The source contract for this chapter is narrow. Claims about implemented behavior come from the repository paths in Table 3.2, not from generic vulnerability management practices.

NVD lookup	apps/api/app/scanning/nvd_client.py, cve_service.py	CPE and keyword lookup, parser behavior, Redis caching, rate backoff
CPE and firmware hints	apps/api/app/scanning/cpe_builder.py, router.py	CPE construction, operating-system CPE exclusion from actionable lookup, local firmware-CVE map
OpenVAS	apps/api/app/scanning/openvas_adapter.py, progressive_vuln_phases.py	Greenbone Management Protocol workflow, report parsing, feed status checks, and background phase merge
Nuclei	apps/api/app/scanning/nuclei_adapter.py, progressive_vuln_phases.py	Command invocation, target construction, JSON-line parsing, partial timeout output

Chapter 3: Vulnerability Assessment

Default credentials	apps/api/app/scanning/default_creds_adapter.py	Opt-in credential checks, vendor-first credential selection, protocol mapping, and false-positive guardrails
Validation	apps/api/app/scanning/cve_validation_service.py, cve_validation_router.py	Validation records, noisy-CVE aggregation, suppressed (cve_id, device_ip) pairs
Tests and lab	apps/api/tests/test_nvd_client.py, test_cve_validation.py, test_progressive_vuln_phases.py, test_openvas*.py, test_default_creds*.py, student-lab/ground-truth.json, student-lab/phase3-lab.md, student-lab/phase3-homework.md	Measured unit behavior, mocked adapter behavior, and simulated lab ground truth

Table 3.2. Source map for Chapter 3 claims.

Threat Model and Assumptions

Chapter 3 assumes a defender operating under partial visibility. The defender knows some hosts, identities, service surfaces, and firmware hints. The defender cannot know the true device state with certainty. Devices may be patched without updating banners. Vendors may reuse product strings across models. Local policy may block active probes. Some devices may expose no local management service at all.

The adversary benefits from those gaps. An attacker can target devices that are poorly identified, excluded from active scans, hidden behind broad product families, or treated as low-consequence appliances. The defender's vulnerability process need not be wrong everywhere for an attacker to benefit. One blind spot is enough.

The chapter also assumes safety constraints. Active vulnerability testing can disrupt fragile devices. Credential testing can trigger lockouts or alerts. Clinical, industrial, and building-control networks may require written permission, maintenance windows, or read-only policies. This chapter does not treat destructive validation as acceptable evidence gathering. When confirmation would be unsafe, the correct result is unresolved risk with a clear reason.

The key assumptions are:

1. No single method is authoritative.
2. A result with no CVE matches is not the same as no risk.
3. A successful default-credential attack is direct evidence, but it still requires proper handling.
4. A high CVSS score is not an automatic first action.
5. A disabled, missing, timed-out, or policy-blocked tool produces an assessment gap, not safety.

3.1 The Chapter Boundary

Chapter 3: Vulnerability Assessment

Vulnerability assessment is the bridge between knowing what a device is and understanding how it might be attacked.

It is not a discovery. Chapter 1 already answered whether a host exists. It is not enrichment. Chapter 2 already explained which identity claim is supported by the evidence. It is not exploitation. Later chapters may model attack paths, run controlled offensive tests, simulate fixes, or plan remediation. Chapter 3 stops earlier. It decides which weaknesses should enter those later stages as evidence-backed claims.

Vulnerability Assessment Pipeline

Identity artifacts remain separate from passive, active, credential, and prioritization evidence.

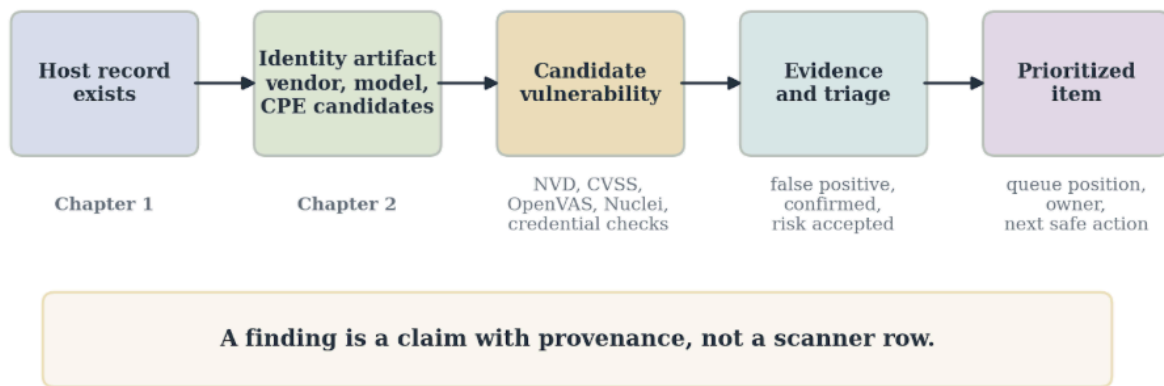


Figure 3.1. Vulnerability assessment pipeline. Chapter 3 starts with Chapter 2 identity artifacts and preserves separate evidence streams before prioritization.

Four terms keep the boundary clean.

Candidate vulnerability	A weakness that may apply based on identity, CPE, banner, template, or scanner evidence	NVD returns a CVE for a wildcard Hikvision camera CPE
Confirmed weakness	A weakness supported by direct or strong corroborating evidence	A Nuclei template verifies a vulnerable endpoint, or default credentials succeed
False positive	A candidate who does not apply to this device after triage	The product family is vulnerable, but this firmware build is patched
Prioritized item	A finding selected for action under operational constraints	A default credential on a camera VLAN moves ahead of a weak medium CVE match

Chapter 3: Vulnerability Assessment

Table 3.3. Chapter 3 vocabulary. The words prevent the tool output from becoming operational truth too quickly.

The biggest pitfall is jumping to conclusions too soon. Seeing an IP, product name, and CVE in a row might tempt you to declare, “the device has CVE-X.” But that shortcut skips the evidence. A stronger statement would be: “The device is a candidate for CVE-X because the enrichment pipeline produced CPE-Y from sources A and B; confidence is high, medium, or low; and active checks did or did not back it up.”

Although this process may appear slower, it serves as a safeguard against costly errors. Vulnerability management frequently fails by either accepting every scanner result, leading to excessive noise, or by disregarding them entirely, overlooking genuine threats. The methodology presented here prioritizes evidence provenance, rigorous evaluation, and transparent documentation of uncertainties.

3.2 From Identity to Applicability

Chapter 2 hands Chapter 3 a host record that may contain services, vendor, model, device type, firmware source, CPE candidates, and confidence. The quality of that record controls the quality of every CVE claim that follows.

A Common Platform Enumeration string is a naming key. It is useful only when it names the right thing at the right granularity. A narrow key can miss vulnerabilities if the vendor’s naming conventions do not match the database’s naming conventions. A broad key can match too many vulnerabilities. A wildcard version can improve recall at the cost of precision. A conflicting CPE set should not be averaged into confidence. It should remain visible.

CPE-to-CVE Lookup Flow

CPE construction, deduplication, cache lookup, and NVD requests produce passive vulnerability evidence.

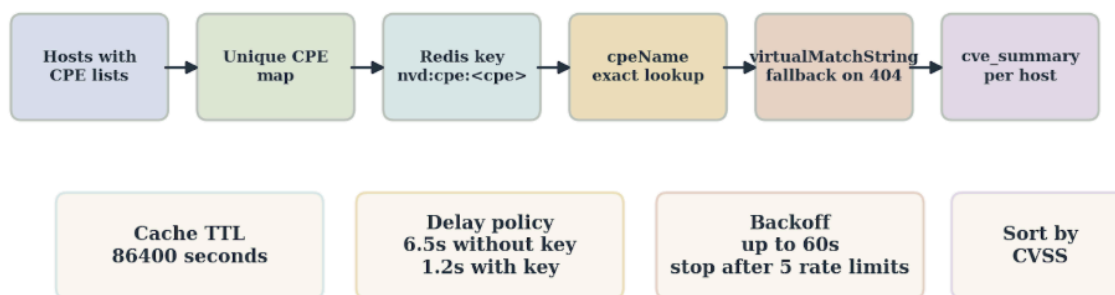


Figure 3.2. CPE-to-CVE lookup flow. CPE construction, deduplication, cache lookup, and NVD requests determine the passive vulnerability evidence available to the analyst.

Breakwater builds CPE strings in `apps/api/app/scanning/cpe_builder.py`. The service path handles conventional service products, HTTP server headers, ONVIF manufacturer and firmware fields, SSDP server and model data, RTSP server headers, and TLS issuer hints. It also contains identity-based CPE mappings for devices that may expose few or no open services, such as cloud-first cameras and media players.

The implementation makes one restraint explicit in `apps/api/app/scanning/router.py`: `nmap` operating-system CPEs are not merged into actionable CVE lookups. The comment explains why. Major-version operating-system CPEs, such as a broad Linux kernel or Android family, can cause

Chapter 3: Vulnerability Assessment

virtualMatchString queries to return large sets of unrelated CVEs. The operating-system CPEs may remain useful for display or later reasoning, but service and identity CPEs drive Chapter 3's CVE lookup.

An increased number of CPEs does not necessarily yield improved results. Broad CPEs may appear comprehensive but can obscure meaningful distinctions. CPE generation should be approached as a deliberate process of evidence transformation, rather than a procedural formality.

Exact product and version	High precision	Misses if NVD uses a different product name	Trust more, but inspect the affected CPE list
Product family without version	Higher recall	Overmatches patched or unrelated builds	Mark as a candidate and seek version evidence
Identity-only wildcard	Useful for sparse devices	Broad and noisy	Keep confidence low until corroborated
Conflicting CPEs	Preserves uncertainty	May mix incompatible product claims	Do not collapse; resolve identity first

Table 3.4. CPE precision and vulnerability applicability.

Before Chapter 3 calls the NVD, the analyst should apply a simple CPE quality check.

First, ask whether the CPE came from a structurally strong source. ONVIF manufacturer, model, and firmware evidence is stronger than a generic HTTP server header for cameras. An mDNS service type is stronger than a broad Google MAC (Media Access Control) Organizationally Unique Identifier (OUI). A service product string like nginx may identify software, but says little about the embedded product that shipped it.

Second, ask whether the version is real. V5.5.52 build 200915 carries more product meaning than a major-only kernel family. A wildcard version may be the best available evidence, but should not be read as identifying a vulnerable build. Breakwater's `_normalise_version()` function in `cpe_builder.py` extracts numeric prefixes because CPE versions need a standard form. That standardization is useful, but it can also discard valuable information that matters to a vendor advisory.

Third, ask whether the CPE names software, firmware, hardware, or an application. `cpe:2.3:a` and `cpe:2.3:o` do not carry the same meaning as `cpe:2.3:h`. A hardware CPE for a camera model may point to firmware vulnerabilities, but the fix may still be a firmware branch rather than a hardware replacement. Students should be able to identify which layer the CPE names belong to before interpreting the CVE result.

Fourth, preserve alternatives. If HTTP suggests one product family and ONVIF another, Chapter 3 should not average the two into a clean answer. It should record the conflict and either seek better identity evidence or lower confidence in any CVE match from the weaker source.

The student lab makes the dependency concrete. In `student-lab/ground-truth.json`, the simulated network contains 26 devices. The Hikvision camera at 172.30.0.10 is listed as DS-2CD2143G2-I with firmware V5.5.52 build 200915, protocols `http/80`, `rtsp/554`, `tcp/8000`, and `ssdp`, default credentials `admin:12345` and `admin:Hikvision`, and known CVEs CVE-2021-36260 and CVE-2017-7921. The Google Chromecast at 172.30.0.20 is a media player with `http/8008`, `tcp/8009`, `tls/8443`, `mdns`, firmware 1.68.382489, and a known CVE entry CVE-2019-2099.

Chapter 3: Vulnerability Assessment

In the lab, students should use these device profiles to practice mapping findings to concrete evidence sources, such as protocol banners, credential testing results, and CVE matches. For each device, students are expected to document the rationale behind associating a specific vulnerability, credential, or protocol with the device, clearly noting the confidence and any supporting evidence. Comparing the two device profiles, students should be able to identify not only the differences in exposure and risk, but also reflect on how evidence paths vary depending on device type, available data, and protocol surface. These actions will help build habits of documenting why a finding is relevant, tracing it to observed facts, and drawing clear distinctions between candidate, confirmed, and prioritized weaknesses.

These two devices exemplify distinct analytical considerations. The Hikvision camera provides robust local evidence and direct credential risk, whereas the Chromecast demonstrates ecosystem-specific evidence and lacks default credentials in the lab environment. Categorizing both simply as 'IoT devices with CVEs' neglects critical distinctions. Applicability is determined by identity, exposure, and the nature of available evidence.

3.3 NVD and CVE Lookup

The National Vulnerability Database lookup path is the passive side of Chapter 3. It asks which CVE records are associated with a CPE or keyword. It does not touch the target device. Its risk is not device disruption. Its risk is analytical overreach.

The client lives in `apps/api/app/scanning/nvd_client.py`. The service layer lives in `apps/api/app/scanning/cve_service.py`. The environment controls are visible in the `.env` file. example: `BREAKWATER_CVE_LOOKUP_ENABLED=true`, `BREAKWATER_NVD_API_KEY=`, and the related NVD configuration in `apps/api/app/scanning/config.py`.

3.3.1 Query Mechanics

The CPE lookup function first tries exact matching through the NVD `cpeName` parameter:

```
cpeName=<full CPE 2.3 string>
```

If the NVD returns HTTP 404, the implementation falls back to `virtualMatchString`. The fallback strips trailing wildcard fields and requests a broader prefix match. This is a deliberate precision-recall shift. Exact `cpeName` is cleaner when it works. `virtualMatchString` is useful when the exact CPE is not registered or when the version granularity is incomplete. It also requires more skepticism.

The parser prefers CVSS version 3.1 metrics, then version 3.0, then version 2. It extracts the CVE ID, English description, base severity, base score, Common Weakness Enumeration (CWE) identifier when present, affected CPE criteria, references, publication time, and last modification time. The parser does not claim exploitability; it produces consistent CVE records.

3.3.2 Caching and Rate Discipline

The National Vulnerability Database is a shared external service. A scanner that calls it once per host per CPE can become slow and noisy. Breakwater reduces that cost in three ways.

First, `lookup_cves_cached_batch()` deduplicates CPEs across hosts. If ten identical devices produce the same CPE, the service performs a single lookup and distributes the result to all ten hosts.

Second, Redis caches each CPE under the key pattern `nvd:cpe:{cpe}`. The default `cve_cache_ttl` in `apps/api/app/scanning/config.py` is 86400 seconds. The cache stores serialized CVE records and converts datetime objects to ISO strings.

Chapter 3: Vulnerability Assessment

Third, uncached NVD calls are sequential and delayed. The implementation waits 6.5 seconds between calls without an API key and 1.2 seconds with one. On rate-limit responses, it uses exponential backoff capped at 60 seconds. After five consecutive rate-limit events, it fills remaining CPEs with empty results rather than stalling the scan.

The NVD path in the progressive pipeline uses this sequence: collect CPEs from each host, deduplicate them across the batch, check Redis for cached results, call the NVD only for misses, then distribute the returned CVEs to matching hosts. The important analytical point is that a cache hit and a live API call have the same applicability limit. Both return database association evidence, not proof that the device is exploitable.

The result attached to a host is `cve_summary`. It contains a total count, counts for critical/high/medium/low, and a capped list of CVE records sorted by base score in descending order. The default cap, `cve_max_per_host`, is 50.

This design draws a key line for students: caching and backoff keep the system running smoothly, but cannot guarantee a CPE match is correct. The real test remains the quality of the identity and how well findings are triaged.

3.3.3 Reading the Returned Record

The returned CVE record should be read as a compact evidence packet.

The `cve_id` is the stable identifier. The description states the public natural-language claim, usually at the product or component level. The severity and base score provide a standardized estimate. The Common Weakness Enumeration field, when present, hints at the weakness category. The affected product list is often the most important part for applicability because it names the CPE criteria recorded in the database. References carry vendor advisories, patches, mitigations, exploit notes, or third-party analysis.

Readers often focus solely on the description field; however, this is insufficient. Descriptions may reference “certain firmware versions” or “affected devices,” but the affected CPE list and vendor advisories establish actual applicability boundaries. When only vendor and device type are available, descriptions may seem persuasive even if the version does not align with the actual version. In contrast, when model and firmware information are present, the affected CPE criteria can be evaluated with greater precision.

Publication and modification times are important but should not be mistaken for indicators of exploitability. A recent modification may reflect a record update rather than a new vulnerability. An old publication date does not always mean low priority; for example, an old camera CVE with a public exploit and exposed management interface can remain critical for years.

The code does not treat reference tags as proof. It extracts them and uses selected tags to generate fix proposals. That is the right level of automation for this stage. A reference-tagged Patch is useful, but the analyst still needs to confirm that the patch applies to the device branch, region, hardware revision, and maintenance process.

3.3.4 Fix Proposals

`cve_service.py` also generates remediation hints. The order matters:

1. NVD references tagged Patch
2. References tagged Vendor Advisory.
3. references tagged Mitigation

Chapter 3: Vulnerability Assessment

4. severity-based generic advice
5. CWE-based hints
6. device-type hints

That order is conservative. A vendor patch is more specific than generic isolation advice. A Common Weakness Enumeration hint may be useful, but it is not proof that the device can be remediated that way. A device-type hint, such as camera isolation, is useful as a control pattern, not as evidence that a particular firmware patch exists.

Remediation proposals serve as preliminary guidance, not definitive change tickets. Formal change tickets require details such as asset ownership, maintenance scheduling, device role, current firmware version, and a rollback strategy.

3.4 CVSS as Structure

The Common Vulnerability Scoring System is frequently reduced to a sortable column of scores. However, an ordered list of scores lacks substantive value if the underlying vector details are not examined.

Version 3.1 base scoring describes exploitability and impact through eight metrics:

AV	Attack Vector
AC	Attack Complexity
PR	Privileges Required
UI	User Interaction
S	Scope
C	Confidentiality impact
I	Integrity impact
A	Availability impact

Table 3.5. CVSS v3.1 base metrics.

Chapter 3: Vulnerability Assessment

CVSS Base Vector

The numeric score is a summary. The vector fields explain what kind of danger the score represents.

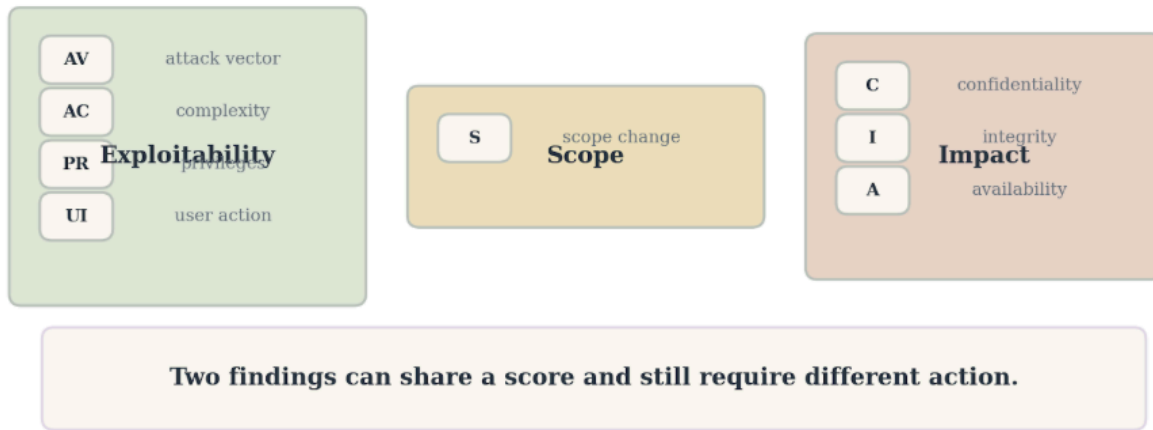


Figure 3.3. CVSS base vector. The numeric score summarizes the vector; the vector explains the kind of danger.

A National Vulnerability Database record may expose the vector as a compact string, such as CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:L/A:N. Students should learn to read that string as a claim about the attack path and impact, not as decoration around the score.

Consider two high-severity findings. One has a network attack vector, low complexity, no privileges required, no user interaction, high confidentiality impact, low integrity impact, and no availability impact. Another has the same exploitability fields but high availability impact and no confidentiality impact. The scores may be similar, but operational priority may differ sharply.

On a media device, confidentiality may matter most if the finding exposes account tokens or private metadata. In a clinical hallway, camera availability may matter if the camera supports safety monitoring. On a router, integrity may dominate because configuration tampering changes the rest of the network. The score helps only when the analyst reads the vector in context.

Breakwater currently sorts CVE summaries by base severity and score, which is useful for initial ordering but not for final prioritization. The chapter's framework later incorporates evidence strength, exposure, device role, and consequence.

The severity bands are familiar:

0.0	None
0.1-3.9	Low
4.0-6.9	Medium

Chapter 3: Vulnerability Assessment

7.0-8.9	High
9.0-10.0	Critical

Table 3.6. CVSS base severity bands.

Severity bands serve as general indicators rather than absolute rules. For example, a critical issue on a laboratory device may be less significant than a high-confidence, high-severity finding on a network-exposed camera. In some cases, a medium-severity vulnerability at a critical network chokepoint warrants more urgent remediation than a high-severity issue on an isolated device. Analysts must always contextualize priorities rather than relying solely on numerical scores.

A database lookup reasons from identity. OpenVAS probes the network.

Breakwater integrates OpenVAS through Greenbone Management Protocol (GMP) in `apps/api/app/scanning/openvas_adapter.py`. It connects to `gvmd` over TLS/TCP on port 9390 by default. The optional dependency is `python-gvm`; if it is missing, `_connect_gmp()` raises a runtime error with an installation hint.

OpenVAS is disabled unless the OpenVAS URL is configured. The derived flag is visible in `apps/api/app/scanning/config.py`: `openvas_enabled` is set from `openvas_url`. The `.env` example file exposes `BREAKWATER_OPENVAS_URL`, `BREAKWATER_OPENVAS_USERNAME`, and `BREAKWATER_OPENVAS_PASSWORD`.

3.5.1 The GMP Lifecycle

The adapter performs a complete GMP workflow:

1. Connect and authenticate.
2. Find a scan config.
3. Find a scanner.
4. Find a port list.
5. Create a target.
6. Create a task.
7. Start the task.
8. Poll until completion or timeout.
9. Fetch the report.
10. Parse results.
11. Delete the task and target.

Chapter 3: Vulnerability Assessment

OpenVAS GMP Lifecycle

The adapter creates targets and tasks through GMP, polls with fresh connections, parses the report, and cleans up.

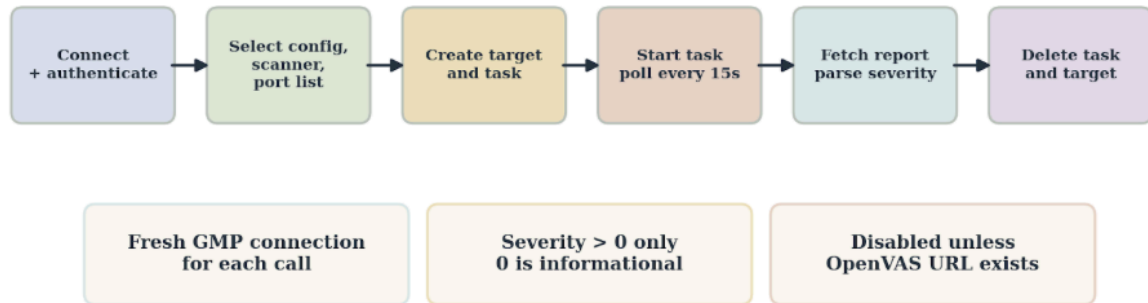


Figure 3.4. OpenVAS GMP lifecycle. The adapter creates targets and tasks, polls with fresh connections, parses the report, and cleans up.

The preferred scan config order in the code is exact: Full and fast, then Base, then Full and very deep. If none are found, the adapter falls back to the first available config. Scanner selection prefers a scanner whose name contains OpenVAS, then falls back to the first scanner. Port-list selection prefers common Greenbone port lists, then falls back to the first available list.

The implementation uses a fresh GMP connection for each polling and report call through `_gmp_call()`. This is not cosmetic. Long-running scans can suffer broken pipes. Opening a new TLS connection, authenticating, running a single operation, and closing it makes the adapter more tolerant of long scan durations and server-side connection churn.

3.5.2 Report Parsing

The parser groups report results by host IP. It skips findings with a severity of 0 or lower. It maps numeric severity into labels:

9.0 or higher	Critical
7.0 to 8.9	High
4.0 to 6.9	Medium
greater than 0 and less than 4.0	Low

Table 3.7. OpenVAS severity normalization in `openvas_adapter.py`.

The parser limits description and solution text to 500 characters to prevent downstream records from becoming unwieldy. However, this truncation may omit critical remediation details or nuanced context.

Chapter 3: Vulnerability Assessment

Consequently, analysts and students should review the complete Greenbone report before drafting final remediation instructions. Exclusive reliance on truncated text risks overlooking essential information; consulting the full report ensures comprehensive guidance and context.

OpenVAS shines by watching how devices behave, catching service-level weaknesses that CPE lookups might miss. Its results depend on probe policy, authentication, feed freshness, and the target's capacity. If active probes skip a device, a clean OpenVAS report does not mean it is safe. Slow or delicate devices might time out, and noisy findings still need careful sorting.

Therefore, Chapter 3 does not present OpenVAS as a definitive scanner. It is an evidence source with a different error profile than the NVD.

3.5.3 Operational Safety

OpenVAS is active. That word should change the analyst's posture.

An active scan sends probes that may be unusual for embedded systems. Many modern devices tolerate them. Some fragile or safety-critical devices should not receive them without policy approval. The same scan profile that is reasonable for a test VLAN may be inappropriate on a clinical network, an industrial controller segment, or a building-management subnet during business operations.

The code provides several safety controls: OpenVAS is disabled unless configured, scan configuration can be overridden via environment variables, timeouts prevent indefinite tasks, and cleanup removes tasks and targets after completion or interruption. Background execution ensures long scans do not block the pipeline. These controls are necessary but do not authorize unrestricted scanning.

One core principle: a host left out of testing is not proven safe—it is simply untested by that method. Final reports should clearly state whether OpenVAS found nothing, never ran, timed out, or was skipped by policy, since each tells a different story. Distinguish between report severity and target consequence. OpenVAS may report medium severity for a device routing traffic between zones, or high severity for an isolated lab camera. The parser assigns severity but does not account for business role; this context is considered during prioritization.

3.6 Nuclei: Template-Based Checks

Nuclei provide another active evidence path in Chapter 3. It runs templates against targets and emits machine-readable findings. Breakwater wraps the local nuclei command-line tool in `apps/api/app/scanning/nuclei_adapter.py`.

Nuclei scanning is off by default. `.env.example` sets `BREAKWATER_NUCLEI_ENABLED=false`. The default tags and severity filters are:

```
BREAKWATER_NUCLEI_TAGS=network,dns,ssl,default-login,cve
```

```
BREAKWATER_NUCLEI_SEVERITY=low,medium,high,critical
```

```
BREAKWATER_NUCLEI_TIMEOUT=600
```

```
BREAKWATER_NUCLEI_RATE_LIMIT=100
```

The adapter checks whether the nuclei binary is present. If not, it returns an empty result and logs the missing binary. When Nuclei is enabled and present, the adapter writes targets to a temporary file and runs the command-line tool with JSONL output, no color, no raw payloads, and no template body. It also applies a per-template timeout of 15, the configured rate limit, a bulk size of 15, and a concurrency of 10.

Chapter 3: Vulnerability Assessment

The implementation path is deliberately explicit: the adapter builds targets from known ports, the command emits JSON lines, and the adapter groups non-informational findings by host. That sequence matters more than template volume because it defines what evidence the pipeline is allowed to claim.

The target construction happens in `apps/api/app/scanning/progressive_vuln_phases.py`. If a host has open ports, the phase creates URL targets for each port. Known TLS ports use `https://`; other ports use `http://`. If a host has no open ports, the phase falls back to the bare IP.

The parser reads JSON lines one at a time. It skips malformed lines. It extracts the IP address from `ip`, then `host`, then `matched-at` when needed. It skips INFO severity because those templates are detection-only in this pipeline. It uses the first CVE ID from the template classification when present; otherwise, it uses the template ID. It stores the name, severity, base score (if parseable), description, matched location, and extracted results.

One implementation detail matters for long scans: the adapter keeps partial results if the global timeout expires. It streams stdout with `select`, kills the process on timeout, drains available output, and then parses what it captured. This behavior is more useful than an all-or-nothing timeout, but the finding set must be read as partial.

Nuclei excels at confirming specific weaknesses, but only when appropriate templates are available. Vulnerabilities lacking a corresponding template may go undetected. Templates can become outdated or fail to identify issues if a device employs unconventional protocol paths. While Nuclei provides strong evidence when verifying actual behavior, its coverage is inherently limited.

Nuclei demonstrates how target construction impacts coverage. Using only a bare IP address can cause tools to default to HTTP or HTTPS. The progressive phase assigns specific targets based on known ports, such as `https://ip:8443` or `http://ip:8008`. This is important for IoT devices, where management and metadata endpoints often use nonstandard ports.

The parser's decision to skip INFO severity is intentional. Technology detection is useful for enrichment, but Chapter 3 does not inflate vulnerability counts with observations like "this server appears to be nginx." Templates that identify only technology belong in Chapter 2; those that verify vulnerable behavior or unsafe configurations belong in Chapter 3.

When Nuclei and database lookup disagree, neither wins automatically. A Nuclei finding with a clear matched endpoint may outrank a weak CPE match. A failed template check does not invalidate a CVE candidate if the template does not cover the exact vulnerable path. The analyst should ask what the template tested, which route it took, which response it matched, and whether the template's assumptions fit the device.

3.7 Default Credentials

When default credential testing succeeds, it is the gold standard for confidence in this chapter. Instead of guessing from product names, it tries logging in with known factory credentials and sees what happens.

That directness is also why it is opt-in. `.env.example` sets:

```
BREAKWATER_DEFAULT_CREDS_ENABLED=false
BREAKWATER_DEFAULT_CREDS_TIMEOUT=5.0
BREAKWATER_DEFAULT_CREDS_CONCURRENCY=3
BREAKWATER_DEFAULT_CREDS_PROTOCOLS=http,ssh,telnet,rtsp,onvif
```

Chapter 3: Vulnerability Assessment

The checker in `apps/api/app/scanning/default_creds_adapter.py` supports HTTP, SSH, Telnet, RTSP, and ONVIF. It maps ports to protocols: 22 to SSH, 23 to Telnet, 80, 443, 8080, 8081, 8443, and 8888 to HTTP, and 554 to RTSP. ONVIF is checked on candidate ports 80, 8080, and 8899 when those ports are open.

Default Credential Assessment

Credential checks are opt-in active authentication attempts, not passive version hints.

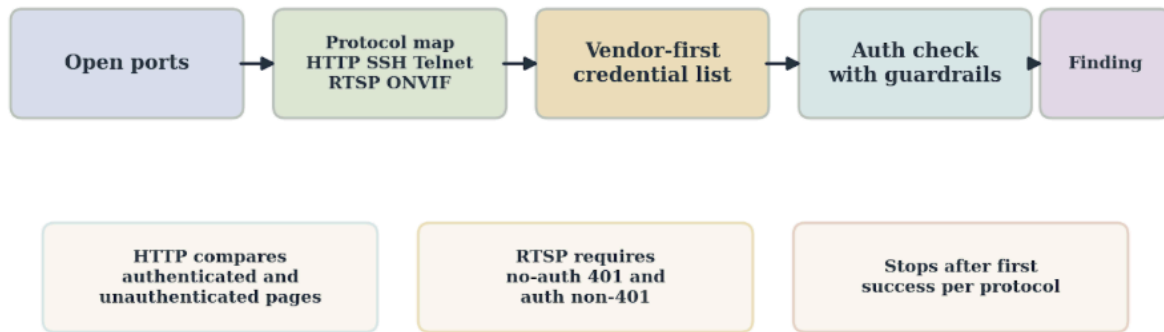


Figure 3.5. Default credential assessment. The checker maps open ports to protocols, tries vendor-specific credentials before generic ones, and uses false-positive guardrails.

Credential selection is vendor-aware. Hikvision HTTP credentials include `admin:12345` and `admin:Hikvision`. Dahua includes `admin:admin` and `admin:.`. Axis includes `root:pass` and `root:.`. Ubiquiti includes `ubnt:ubnt`. The generic set is tried after vendor-specific entries, and duplicate pairs are removed.

HTTP does not report success merely because the response is not 401. It treats 401 and 403 as failures, 429 as redirects, and server errors as inconclusive, and only 2xx responses as candidates. It then sends an unauthenticated request. If the unauthenticated response is also successful with a similar body length, the checker treats the page as public rather than authenticated.

RTSP uses a similar comparison. The no-auth request must return 401, and the authenticated request must return a non-401 status code. If no-auth also succeeds, the endpoint may be open, but it is not counted as a default credential finding.

The checker terminates after identifying the first successful credential for each protocol and host, pausing for 0.5 seconds between attempts. Passwords are not logged. The result includes the credential pair, necessitating careful handling in downstream storage and reporting. A successful default-credential finding requires immediate remediation, such as rotating credentials, turning off risky protocols, segmenting the device, or restricting remote management. Exploiting weaknesses is never appropriate; remediation must always adhere to established safety protocols.

Default credential testing introduces ethical considerations. While a CVE lookup queries an external database, credential checks send authentication attempts to devices. Even with public factory defaults, these tests can generate logs, trigger lockouts, or violate customer rules of engagement. For this reason, the phase is turned off by default.

The lab is intentionally safer than a production network. `ground-truth.json` lists default credentials so that students can reason about expected findings. In production, the scanner should not assume a vendor default list is acceptable to try merely because a port is open. The operator needs scope, authorization, rate limits,

Chapter 3: Vulnerability Assessment

and a plan for handling the secret. Distinguishing between the controlled laboratory environment and the ethical and legal complexities of real-world operations is essential. Actions deemed safe or permissible in lab exercises may have significant consequences if performed in production environments without explicit authorization. Students should not transfer laboratory practices directly to operational systems without evaluating legal, policy, and ethical considerations. Professional responsibility requires confirming authorization, understanding organizational policies, and assessing the potential impact of each scanning action before execution.

A successful credential test confirms that the tested protocol accepted the credential, but does not prove all interfaces use the same credential, that firmware CVEs apply, or that the device is exploitable from every network segment. However, it does justify immediate governance action due to failed authentication at a basic boundary.

3.8 False Positives and Validation

A false positive is not just a result you wish were not there. It is a candidate that does not hold up when checked for true applicability.

In IoT and OT vulnerability assessment, false positives often come from five sources.

CPE imprecision	The CPE names a family broader than the device	Does the affected CPE list include this model and version?
Backported patch	Banner stays old while the fix is present	Does vendor firmware history show a backport?
Stale banner	Service string does not reflect installed code	Can another source confirm the version?
Configuration dependency	CVE applies only when a feature is enabled	Is the vulnerable service or setting present?
Active-scan interpretation	Probe response is misread	Can a safer second check reproduce the behavior?

Table 3.8. Common false-positive mechanisms.

Breakwater records analyst decisions in `apps/api/app/scanning/cve_validation_service.py`. A validation can be one of three types: false positive, confirmed, or risk accepted. The service can create, list, delete, aggregate noisy CVEs, compute validation stats, and return suppressed CVE/device pairs.

Chapter 3: Vulnerability Assessment

False-Positive Validation Profile

Analyst decisions suppress only the affected CVE and device pair; broad global suppression hides risk.

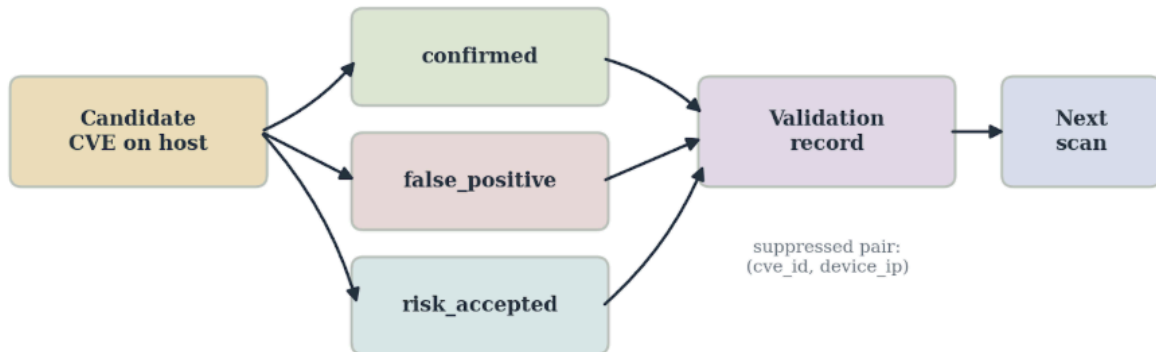


Figure 3.6. Illustrative false-positive sources. This figure is a teaching aid, not a measurement from the current lab; the chapter’s implemented validation behavior is the pair-level suppression mechanism in `cve_validation_service.py`.

Suppression is deliberately tied to each `(cve_id, device_ip)` pair. This matters because a CVE might be a false alarm on one device but a real threat on another. Devices can be patched, rolled back, or replaced, changing status over time. Suppressing findings globally risks hiding real dangers, but pair-level suppression keeps noise down without losing sight of the evidence.

The progressive pipeline marks where suppressed pairs should be loaded and passed into `lookup_cves_cached_batch()`. The service layer accepts the `suppressed_pairs` argument and skips matching pairs during CVE distribution. That distinction is important for code-grounded writing. The suppression mechanism exists in the service; full pipeline loading of organization-specific pairs is a handoff point, not something the chapter should overclaim as universal runtime behavior.

Good validation notes name the reason. “False positive” is weak. “CVE applies to firmware before V5.5.800; ONVIF reports V5.7.1 and vendor advisory confirms patch branch” is useful. The note lets future analysts challenge, reuse, or retire the decision when the device changes.

3.9 Correlation

Correlation is about comparing different streams of evidence to see if they independently back up the same claim—not just tallying the same finding over and over.

NVD, OpenVAS, Nuclei, and default credential testing observe different things:

NVD	CPE-to-CVE association	Broad known-vulnerability coverage	Depends on identity and version quality
OpenVAS	Network probe and Network Vulnerability Test (NVT) result	Mature active assessment	Slow, policy-dependent, feed-dependent

Chapter 3: Vulnerability Assessment

Nuclei	Template-driven target behavior	Targeted active checks	Template coverage and staleness
Default credentials	Successful authentication	High-confidence direct evidence	Narrow vulnerability class and active risk

Table 3.9. Evidence paths by method.

Correlation Without Double Counting

Agreement strengthens evidence only when the methods observe different things.

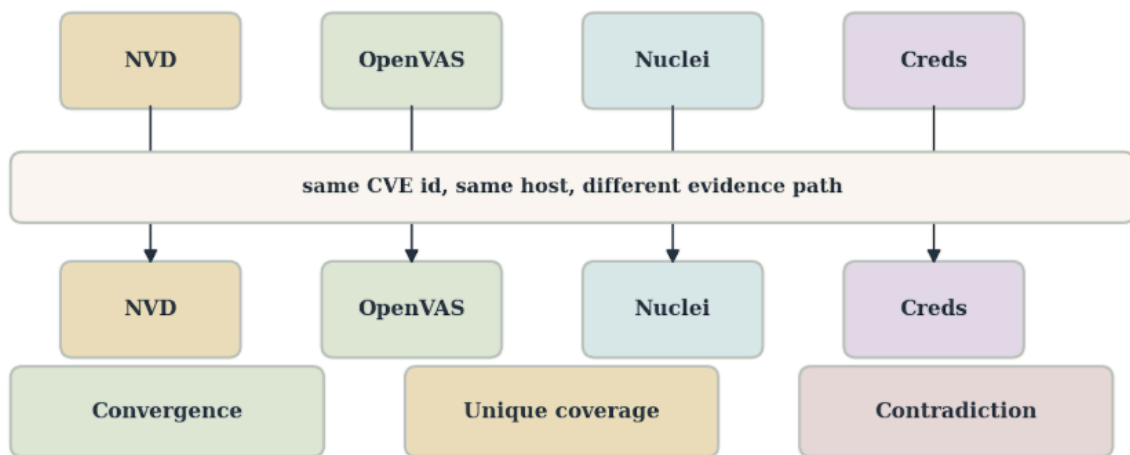


Figure 3.7. Correlation without double counting. An agreement strengthens evidence only when methods observe different things.

Three correlation patterns matter.

Convergence occurs when different methods support the same claim. A CPE match indicates that a camera's firmware is associated with a CVE, and an active template verifies the affected endpoint. The evidence is stronger because the methods depend on different observations.

Unique coverage occurs when one method finds something the others are not designed to find. A default credential success may have no CVE ID. That does not make it less important. It means it belongs to a different evidence class.

Contradiction occurs when methods disagree. A CPE match says vulnerable, but the active probe fails. That could mean the device is patched, the template is incomplete, the service is filtered, the scan lacked credentials, or the vulnerability requires a configuration that the device does not expose. Contradictions should survive in the analyst's notes until resolved.

A simple confidence ladder is useful in class:

Chapter 3: Vulnerability Assessment

Successful authentication finding	High
Active behavior check plus matching identity	High
OpenVAS or Nuclei single-tool finding with clear target behavior	Medium to high
Exact CPE and version match without active confirmation	Medium
Wildcard CPE match without version	Low to medium
Broad family match with conflicting identity evidence	Low

Table 3.10. Evidence-strength ladder for vulnerability findings. The ladder is a teaching tool, not a replacement for analyst judgment.

Active evidence does not always take precedence. Some active checks are limited in scope, whereas a passive vendor advisory may provide highly accurate information. The essential skill lies in evaluating the observations and claims supported by each method.

Three short patterns help students practice correlation.

Pattern one: passive candidate, active confirmation. A Hikvision CPE is vulnerable to CVE-2021-36260, and a selected active check verifies the presence of the vulnerable endpoint behavior. The claim strengthens because the CPE and the active probe depend on different evidence. The queue should treat this as more urgent than a family-level CPE match alone.

Pattern two: direct credential finding without CVE. A device accepts admin:admin, but the NVD query returns no matching CVE. The right conclusion is “no vulnerability.” The right conclusion is “known CVE coverage is weak or absent, but authentication failure is confirmed.” That finding can move to remediation even without a CVE ID.

Pattern three: broad CPE match, no active corroboration. A wildcard CPE produces several high CVEs, but neither OpenVAS nor Nuclei confirms them. The finding should remain in the queue as candidate risk, especially if the device is exposed. Still, the next action may be identity refinement or safe version validation rather than immediate patching. The contradiction is useful because it tells the analyst what to check next.

Correlation notes should be documented within each finding, not kept only as internal analyst knowledge. This allows future reviewers to determine whether a single weak source, a strong source, or multiple independent sources support a queue item.

3.10 Prioritization

A remediation queue is more than just a list sorted by severity.

Severity provides important structure, but priority also depends on whether the finding applies, how strong the evidence is, whether the device is reachable, what role the device plays, and what consequence follows if the weakness is exploited.

Chapter 3: Vulnerability Assessment

Prioritization Inputs

The remediation queue is a judgment over severity, evidence, exposure, role, and consequence.

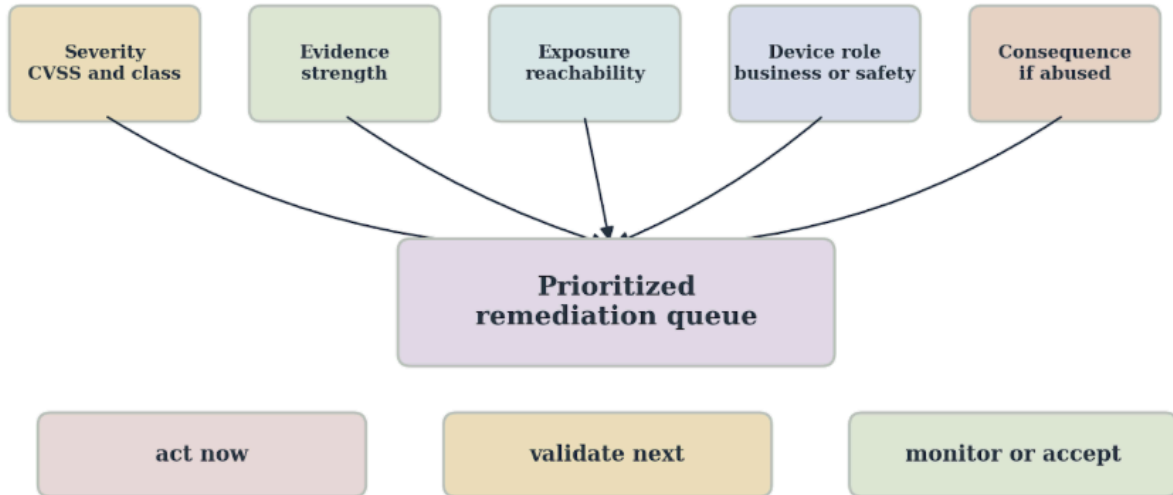


Figure 3.8. Prioritization inputs. The remediation queue combines severity, evidence strength, exposure, device role, and consequence.

The following framework is simple enough for student use and strict enough for operational discussion:

Severity	How bad is the weakness if it applies?
Evidence strength	How directly does the evidence show applicability?
Exposure	Can an attacker reach the affected surface?
Device role	What does the device control, observe, route, or protect?
Consequence	What happens if the weakness is abused?

Table 3.11. Remediation queue inputs.

The queue should produce action bands, not only ranks:

Act now	High consequence and strong evidence	Patch, isolate, rotate credentials, or turn off exposure
Validate next	Severe but uncertain	Gather safe confirmation evidence
Schedule	Real finding with lower urgency	Add to maintenance plan
Monitor	Weak or external dependency	Track until more evidence exists

Chapter 3: Vulnerability Assessment

Accept temporarily	Known risk accepted under policy	Record owner, reason, and review date
--------------------	----------------------------------	---------------------------------------

Table 3.12. Priority bands.

The following examples show how the queue changes when evidence and consequence are included.

Camera accepts the default RTSP credentials	Critical by policy	Direct authentication success	Camera VLAN reachable from the user subnet	Act now
The router has an exact high CVE and exposed management	High	Exact CPE plus open management port	Routes between zones	Act now or validate immediately
Media player has a broad family of CVEs	High	Wildcard CPE only	Isolated entertainment VLAN	Validate next or schedule
Printer has medium CVE	Medium	OpenVAS active finding	Used for sensitive clinical forms	A schedule higher than the score alone suggests
Cloud-only camera has no CVEs	None reported	Sparse local evidence	No open local services	Monitor the identity gap

Table 3.13. Priority examples. The table is illustrative; it shows a reasoning structure. An assessment should not be marked as 'green' if any high-confidence, high-severity, or critical findings remain unresolved. The 'green' status indicates that no significant actions are pending. If risk has been accepted, the status should be recorded as 'accepted.' If remediation is pending a maintenance window, the status should be 'scheduled.' If active confirmation is unsafe, the status should be 'unresolved.' Status terminology must accurately reflect the operational condition. should preserve the truth of the operational condition.

The remediation queue should articulate the next safe step in clear, actionable terms. For example, rather than stating 'patch camera,' provide specific instructions such as: 'Verify the model and firmware via ONVIF or the vendor interface, update default credentials, restrict HTTP/RTSP management to the administrative VLAN, and then schedule a firmware update.' Effective queue items assist asset owners by avoiding technical jargon. The device at 172.30.0.10 serves as a strong teaching example in Chapter 3, as it presents identity, CVE, and credential evidence concurrently. Chapter 2 establishes the identity trace, while Chapter 3 determines which weakness claims can be substantiated for that identity, without conflating the lab answer key with scanner output.

Chapter 3: Vulnerability Assessment

Identity handoff	Hikvision camera, model, and firmware from lab ground truth	Strong enough to construct candidate CPEs, but still distinct from live scanner proof
NVD lookup	CVE-2021-36260 and CVE-2017-7921 are listed in the lab ground truth	Candidate applicability must still be checked against the affected criteria and the firmware branch
Active tools	OpenVAS and Nuclei can add findings when enabled and reachable	Absence of findings is meaningful only when the phase has run and been completed
Credential test	Lab ground truth includes admin:12345 and admin:Hikvision	A successful opt-in login is direct evidence and can outrank a broader CVE hypothesis

Table 3.14. *Hikvision trace compressed to Chapter 3 evidence decisions. The lab ground truth is a grading oracle, not an analytical shortcut.*

The final priority is not “two CVEs plus credentials equals score X.” It is an evidence statement:

172.30.0.10 is a Hikvision camera with exposed HTTP and RTSP services.

The lab ground truth lists CVE-2021-36260 and CVE-2017-7921 for this device.

Default credentials are also present in the lab profile.

If the credential check succeeds in a live run, the finding becomes high-confidence.

The first action is to remove default access and restrict management reachability.

Firmware applicability should be validated next.

That statement is defensible because it distinguishes between evidence and uncertainty.

This device further demonstrates the importance of staged remediation. The initial step should be removing default credentials, as this is directly verifiable and typically low-risk if a backup plan exists. Subsequently, management exposure should be limited, since even patched devices should not have unrestricted administrative interfaces. Firmware validation should follow, potentially requiring vendor-specific information. Proceeding directly from CVE enumeration to firmware updates may overlook the most immediate corrective action: securing default access.

It is important to distinguish between laboratory ground truth and scan-derived evidence. The lab file explicitly lists CVEs and credentials for the simulated device, while real-world scans must infer this information without direct access. The lab serves as a grading oracle, not an analytical shortcut.

Chapter 3: Vulnerability Assessment

3.12 Worked Contrast: Chromecast

The Chromecast at 172.30.0.20 teaches a different lesson. The lab ground truth lists it as a Google media player with http/8008, tcp/8009, tls/8443, mdns, firmware 1.68.382489, no default credentials, and known CVE CVE-2019-2099.

This device is not a camera. Its evidence comes from Cast-specific services, HTTP setup metadata, TLS behavior, and mDNS. A default credential workflow should not dominate its assessment. A camera-oriented CPE should not be applied because the MAC OUI indicates that Google and another Google device in the fleet were cameras. The vulnerability story must follow the identity story.

The Chromecast trace also shows why zero findings can be ambiguous. If database lookup returns no CVEs, that may mean the CPE was too weak, the relevant product name did not match the database, the firmware was not represented, or the device truly had no matching records. If Nuclei returns no findings, it may mean that no applicable templates ran or that the endpoints were unreachable. If default credentials return none, that is expected for this class in the lab, not to prove the host has no risk.

Before planning remediation, compare the Hikvision and Chromecast cases side by side. The Hikvision camera demands quick action on credentials and management controls, while the Chromecast calls for attention to media exposure, metadata leaks, and firmware fit. Even though both are in the same assessment phase, their priorities are worlds apart.

This contrast is valuable in oral exams. Ask students why the camera's default credential should take priority over a weak Chromecast CVE match, even if the numeric severity suggests otherwise. Then ask when the Chromecast should take priority over the camera. Answers should address exposure and consequence. For example, a Chromecast on an isolated home media VLAN differs from one on a production conference network with guest access and sensitive casting history.

3.13 Pipeline Integration

In the progressive scan implementation, CVE lookup is part of the main post-identification vulnerability phase. OpenVAS, Nuclei, and default-credential checks are later phases of vulnerability management. The separation matters because the phases have different runtime costs and safety profiles.

The relevant host fields include:

cve_summary	CVE lookup through <code>lookup_cves_cached_batch()</code>
openvas_vulns	OpenVAS phase
nuclei_vulns	Nuclei phase
default_creds	Default credential phase
firmware_cves	Firmware version lookup in the router and the firmware phase

Table 3.15. Vulnerability-related host fields used by the current codebase.

Chapter 3: Vulnerability Assessment

The `firmware_cves` field is included because it co-resides on the host record. The deeper firmware phase belongs to a later chapter; Chapter 3 uses the field only as additional evidence of a vulnerability when it is already present.

OpenVAS can run in the background. The progressive code includes a background path, so long active scans do not block the main pipeline. Nuclei and default credential phases run against host data collected from Redis or the headless state. The pipeline merges those findings back into the host record and emits progress events.

Users should recognize that scan results may be incomplete at any time. For example, a host may show CVE results before OpenVAS scanning finishes, or have Nuclei findings but lack default-credential results if that phase is disabled. Final reports must clearly state which assessment phases were executed. Reporting vulnerability counts without phase provenance is not defensible. The code emits error events when a vulnerability phase fails. That is a state to preserve, not a reason to produce a clean report in silence. A production report should distinguish at least four states for each method:

Completed	Method ran and produced findings, or no findings
Disabled	Configuration turned the method off
Unavailable	A dependency, binary, service, or credential was missing
Failed or partial	Method started but errored, timed out, or returned partial data

Table 3.16. Method-completion states that it should travel with vulnerability results.

This is especially important for OpenVAS and Nuclei. If Nuclei is disabled by default and the binary is missing, the absence of Nuclei findings says nothing about the target. If OpenVAS is not configured, the absence of OpenVAS findings means no OpenVAS evidence. If default credential checks are disabled, the absence of credential findings means policy avoided active authentication attempts. A mature report states this plainly.

Adversarial Cases

Chapter 3 cannot prove the absence of vulnerabilities.

A device with no CVEs is not necessarily safe. It may have weak identity evidence, missing CPEs, cloud-only behavior, private vulnerabilities, configuration weaknesses, supply-chain issues, or unassigned vulnerabilities. It may also be safe with respect to known public CVEs. The assessment must clarify which interpretation the evidence supports.

The current codebase defends against several practical failure modes:

NVD overuse	Redis cache, cross-host CPE deduplication, delays, backoff
NVD 404 on exact CPE	virtualMatchString fallback
Informational OpenVAS noise	Severity 0 and lower skipped

Chapter 3: Vulnerability Assessment

Long OpenVAS connections	Fresh GMP connection per operation
Missing Nuclei binary	Clean skip with empty result
Nuclei timeout	Partial stdout retained
HTTP credential false positives	An authenticated response compared to an unauthenticated response
RTSP no-auth endpoints	No-auth success is not counted as credential success
Repeated known false positives	Pair-level CVE/device validation support

Table 3.17. Code-level defenses are present in the current implementation.

The codebase also has limits:

Cloud-only devices	Few local services mean sparse active evidence
MAC randomization or spoofing	Vendor hints can mislead identity and CPE generation
Backported patches	Version banners may overstate vulnerability
Template gaps	Nuclei cannot find any selected template tests
OpenVAS policy exclusions	Excluded devices can appear clean
Authentication gaps	Authenticated vulnerabilities may be invisible without credentials
Account lockout risk	Credential checks must remain opt-in and rate-limited
Suppression loading	Service supports suppressed pairs; pipeline integration must load them correctly

Table 3.18. Known limits and open implementation edges.

Adversaries may exploit these limitations. For example, a device may display a deceptive banner, a proxy may cause multiple products to appear identical, or a vulnerable feature may be turned off during scanning. Management interfaces may only be accessible from segments beyond the scanner's reach, and devices may

Chapter 3: Vulnerability Assessment

alter responses to known probes. The appropriate response is to maintain transparent evidence provenance and avoid claims beyond established measurement boundaries.

For students, the adversarial cases are not only attacker stories. They are tests of analytic humility.

If a device spoofs a vendor banner, CPE construction may move. For students, adversarial cases are more than just tales of attackers—they are lessons in analytic humility. JARM) Grouping from Chapter 2 may suggest a relationship that is not product identity. If a vendor backports a patch without changing the exposed version string, NVD may overstate risk. If a scanner lacks credentials, authenticated vulnerabilities may remain invisible. If a firewall allows probes from the scanner but not from realistic attacker positions, reachability may be overstated.

Each scenario requires a nuanced response, not simple acceptance or dismissal. Analysts should verify identities independently, examine affected CPE criteria, compare active and passive findings, document unresolved contradictions, or define safe validation tasks. Conclusions should clearly state any remaining uncertainties.

3.15 What Chapter 3 Hands Forward

Chapter 4 should receive weighted vulnerability claims, not a raw spreadsheet.

A good Chapter 3 handoff contains:

1. host identity and confidence from Chapter 2
2. CPE candidates and their precision
3. CVE candidates with source and affected-product evidence
4. CVSS base score and vector when available
5. active findings from OpenVAS and Nuclei with method provenance
6. default credential findings as direct authentication evidence
7. validation status: false positive, confirmed, risk accepted, or unresolved
8. priority band and rationale
9. explicit unresolved contradictions

This handoff supports attack-graph reasoning. An edge based on a high-confidence default credential differs from one based on a broad wildcard CPE. A path through a safety-critical device is not equivalent to one through a media player. A vulnerability with direct active confirmation is not the same as one inferred from a stale banner.

Subsequent chapters can only perform scoring, simulation, testing, or remediation if Chapter 3 maintains these critical distinctions.

3.16 Student Lab Workflow

The Phase 3 student lab is built to foster disciplined analysis, not just button-pushing.

Chapter 3: Vulnerability Assessment

The lab scope is intentionally limited. `student-lab/phase3-lab.md` uses scan results, CVE lookup, CVE detail, and CVE search endpoints, focusing on interpretation rather than full API exploration. Students can complete the core work with these requests:

```
GET /v1/scanning/smart-scan/{scan_id}/results
```

```
POST /v1/cves/lookup
```

```
GET /v1/cves/{cve_id}
```

```
GET /v1/cves/search?keyword={kw}&limit={n}
```

The lab-facing lookup endpoint is not the only CVE path in the system. The progressive scanner uses `lookup_cves_cached_batch()` to distribute CPE-based results to hosts via Redis, with caching and deduplication. The lab endpoint allows students to adjust host descriptions and observe how results change.

The lab begins with pre-commitment. Before reviewing vulnerability outputs, students select hosts, state their expectations, identify evidence that would challenge those expectations, and decide which method should be strongest for each host. This approach prevents bias from vulnerability reports and ensures honest comparison.

The next exercise compares broad versus narrow CPE strategies. The goal is not to chase the highest CVE count, but to see how a broader description catches more but blurs precision, while a narrower one sharpens accuracy but risks missing matches. Students should walk away able to explain this tradeoff, not just equate more CVEs with better analysis. Students should be able to compare vectors, not only scores. They must explain how attack vectors, required privileges, user interaction, and impact fields change in priority across different operational contexts. The same high score may indicate data exposure on one device and unavailability on another. The numeric score starts the conversation; it does not end it.

The multi-tool exercise then asks students to compare database lookups, OpenVAS, Nuclei, and default credentials on a single host. They should identify convergence, unique coverage, and contradiction. A strong answer distinguishes “database lookup inferred from identity,” “Nuclei matched behavior,” “OpenVAS reported a network test,” and “credential testing authenticated.” A weak answer just counts findings.

The zero-result exercise is equally important. Students find a host with few or no CVEs and explain why that result is ambiguous. Sparse exposure, weak identity, missing CPEs, database coverage gaps, disabled active tools, or a truly low known vulnerability surface can all result in small counts—the lab rewards. Final laboratory outputs should be presented in concise analyst memoranda. Each memo must distinguish between the tool output, the rationale for the output, the evidence strength, and the recommended action. If these four elements are not clearly separated, the analysis is not sufficiently prepared for progression to Chapter 4. If those four layers can be reduced to one sentence, the answer is not ready for Chapter 4.

3.17 Summary

Vulnerability assessment is a disciplined process of determining which weaknesses are genuinely applicable. It begins by establishing device identity, translating it into potential vulnerabilities, confirming what can be safely validated, and then selecting subsequent actions by considering severity, evidence strength, exposure, device role, and potential consequences.

Database lookups cast a wide net for known vulnerabilities, but their reach depends on the quality of CPE data. Severity scores help organize findings, yet the real story lies in the details of the vector. OpenVAS and Nuclei add active, hands-on evidence, each with its own limits and safety checks. Default credential testing stands out for its directness and high confidence. Validating false positives keeps the noise down while

Chapter 3: Vulnerability Assessment

preserving the evidence trail. Correlation only strengthens findings when you recognize how each method differs. Prioritization turns all this evidence into clear, actionable steps by factoring in exposure, device role, and what is truly at stake.

The chapter's guiding rule is straightforward:

A scanner's finding should never be accepted as the final determination.

Instead, it is necessary to articulate the significance of the finding, assess the strength of the supporting evidence, and specify the recommended next steps.

That approach is the bridge between raw tool output and true cyber analysis.

Review Questions

1. Why is a CPE-to-CVE match a candidate vulnerability rather than an automatically confirmed weakness?
2. Give one example where a broad CPE improves recall but harms precision.
3. Why does Breakwater check Redis before calling the NVD API?
4. What does virtualMatchString change about the NVD lookup compared with exact cpeName?
5. Why can two CVEs with similar CVSS scores require different operational priorities?
6. What does OpenVAS observe that NVD lookup does not?
7. What does Nuclei observe that OpenVAS or NVD may miss?
8. Why is a successful default credential check usually stronger evidence than a passive CVE match?
9. Why should false-positive suppression be scoped to (cve_id, device_ip) rather than globally to the CVE?
10. In the Hikvision lab trace, which evidence would you act on first and why?
11. Why can zero CVE results fail to prove that a device is safe?
12. What should Chapter 3 hand to an attack graph in Chapter 4?

Advanced Discussion Prompts

1. A camera has a critical NVD match from a wildcard CPE, but Nuclei and OpenVAS return no active confirmation. Where should it sit in the remediation queue?
2. A media player has one medium CVE but sits on a flat network next to clinical workstations. Should priority rise? What evidence would you need?
3. A default credential works for RTSP, but HTTP management rejects it. What can and cannot be concluded?
4. An analyst marks a CVE as a false positive because "the scanner is noisy." What would make that validation record acceptable?

Chapter 3: Vulnerability Assessment

5. Should active scans run against safety-critical devices by default? What policies control that decision?

References

- National Institute of Standards and Technology, National Vulnerability Database API 2.0 documentation.
- FIRST, Common Vulnerability Scoring System v3.1 specification.
- MITRE, Common Vulnerabilities and Exposures program documentation.
- MITRE, Common Weakness Enumeration documentation.
- Greenbone Community Edition and Greenbone Management Protocol documentation.
- ProjectDiscovery, Nuclei, and Nuclei Templates documentation.
- NIST Special Publication 800-40, Guide to Enterprise Patch Management Planning.
- OWASP, Internet of Things guidance, and authentication weakness categories.

Cross-References

- **Chapter 1:** discovery and host evidence provenance.
- **Chapter 2:** enrichment, fingerprinting, CPE construction, and identity confidence.
- **Chapter 4:** attack graphs and Breakwater Risk Score inputs.
- **Chapter 5:** controlled autonomous testing, where confirmed findings may become test candidates.
- **Chapter 6:** digital-twin simulation of remediation effects.