

Development of an AI Agent Enrichment Workflow at the National Vulnerability Database

Harold Booth and Derek Sappington
National Institute of Standards and Technology

Development of an AI Agent Enrichment Workflow at the National Vulnerability Database

Certain commercial equipment, instruments, or materials, commercial or non-commercial, are identified in this presentation in order to specify the experimental procedure adequately. Such identification does not imply recommendation or endorsement of any product or service by NIST, nor does it imply that the materials or equipment identified are necessarily the best available for the purpose.

The opinions, recommendations, findings, and conclusions in this presentation do not necessarily reflect the views or policies of NIST or the United States Government.

Setting The Table

Vulnerability
Ecosystem


NVD

V-Etalon

Lessons
Learned

Future

Becoming a CVE Record



How is a CVE Used?

Coordination

Public
Disclosure

Identification

Prioritization

Initial CVE Creation



Image: Generated by Google's Gemini 3 (Thinking) April 14, 2026

Industrialization of CVE



We want this

NIST






Image: Generated by Google's Gemini 3 (Thinking) September 17, 2026


And Avoid This



CVE Production Trends




CVE Production Trends



Bug Fixes for Chrome

Bug fixes by milestone



Bug Fixes for Mozilla Firefox

Firefox Security Bug Fixes by Month

All Sources • All Severities



Request for Information (RFI) on Modernizing the National Vulnerability Database in the Age of Artificial Intelligence



Vulnerability Management Process

- Dissemination
- Risk Management
- Remediation

Vulnerability Data and Standards

Development Processes

Vision for the NVD



Released August 12, 2026



Public Comment Period
until October 13, 2026

National Vulnerability Database

 An official website of the United States government [Here's how you know](#) ▼

NIST

≡ NVD MENU

[Information Technology Laboratory](#)

NATIONAL VULNERABILITY DATABASE

NIST NATIONAL VULNERABILITY
DATABASE
NVD

General

+

Vulnerabilities

+

Vulnerability Metrics

+

Products

+

Developers

+

Contact NVD

Other Sites

+

Search

+



[Communications Page](#)



[CVSS v4.0 Support](#)



[2.0 APIs](#)

The NVD is the U.S. government repository of standards based vulnerability management data represented using the Security Content Automation Protocol (SCAP). This data enables automation of vulnerability management, security measurement, and compliance. The NVD includes databases of security checklist references, security-related software flaws, product names, and impact metrics.

NVD Enrichment

Severity Scoring

*Common Vulnerability Scoring System
(CVSS)*

Vulnerability Type

*Common Weakness Enumeration
(CWE)*

Product Applicability

Common Platform Enumeration (CPE)

Informative References

Reference Link Metadata

VULNERABILITIES

🌸 CVE-2021-44228 Detail

Description

Apache Log4j2 2.0-beta9 through 2.15.0 (excluding security releases 2.12.2, 2.12.3, and 2.3.1) JNDI features used in configuration, log messages, and parameters do not protect against attacker controlled LDAP and other JNDI related endpoints. An attacker who can control log messages or log message parameters can execute arbitrary code loaded from LDAP servers when message lookup substitution is enabled. From log4j 2.15.0, this behavior has been disabled by default. From version 2.16.0 (along with 2.12.2, 2.12.3, and 2.3.1), this functionality has been completely removed. Note that this vulnerability is specific to log4j-core and does not affect log4net, log4cxx, or other Apache Logging Services projects.

Metrics

CVSS Version 4.0

CVSS Version 3.x

CVSS Version 2.0

SSVC

NVD enrichment efforts reference publicly available information to associate vector strings. CVSS information contributed by other sources is also displayed.

CVSS 3.x Severity and Vector Strings:



NIST: NVD

Base Score: 10 CRITICAL

Vector: CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:C/C:H/I:H/A:H

ADP: CISA-ADP

Base Score: 10 CRITICAL

Vector: CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:C/C:H/I:H/A:H

QUICK INFO

CVE Dictionary Entry:

[CVE-2021-44228](#)

NVD Published Date:

Dec 10, 2021

NVD Last Modified:

Aug 11, 2026

Source:

Apache Software Foundation

References to Advisories, Solutions, and Tools

By selecting these links, you will be leaving NIST webspace. We have provided these links to other web sites because they may have information that would be of interest to you. No inferences should be drawn on account of other sites being referenced, or not, from this page. There may be other web sites that are more appropriate for your purpose. NIST does not necessarily endorse the views expressed, or concur with the facts presented on these sites. Further, NIST does not endorse any commercial products that may be mentioned on these sites. Please address comments about this page to nvd@nist.gov.

URL	Source(s)	Tag(s)
http://packetstormsecurity.com/files/165225/Apache-Log4j2-2.14.1-Remote-Code-Execution.html 	Apache Software Foundation, CVE	Third Party Advisory VDB Entry
http://packetstormsecurity.com/files/165260/VMware-Security-Advisory-2021-0028.html 	Apache Software Foundation, CVE	Third Party Advisory VDB Entry
http://packetstormsecurity.com/files/165261/Apache-Log4j2-2.14.1-Information-Disclosure.html 	Apache Software Foundation, CVE	Exploit Third Party Advisory VDB Entry
http://packetstormsecurity.com/files/165270/Apache-Log4j2-2.14.1-Remote-Code-Execution.html 	Apache Software Foundation, CVE	Exploit Third Party Advisory VDB Entry
http://packetstormsecurity.com/files/165281/Log4j2-Log4Shell-Regexes.html 	Apache Software Foundation, CVE	Third Party Advisory VDB Entry
http://packetstormsecurity.com/files/165282/Log4j-Payload-Generator.html 	Apache Software Foundation, CVE	Third Party Advisory VDB Entry
http://packetstormsecurity.com/files/165306/L4sh-Log4j-Remote-Code-Execution.html 	Apache Software Foundation, CVE	Third Party Advisory VDB Entry

This CVE is in CISA's Known Exploited Vulnerabilities Catalog



Reference [CISA's BOD 22-01](#) and [Known Exploited Vulnerabilities Catalog](#) for further guidance and requirements.

Vulnerability Name	Date Added	Due Date	Required Action
Apache Log4j2 Remote Code Execution Vulnerability	December 10, 2021	December 24, 2021	For all affected software assets for which updates exist, the only acceptable remediation actions are: 1) Apply updates; OR 2) remove affected assets from agency networks. Temporary mitigations using one of the measures provided at https://www.cisa.gov/uscert/ed-22-02-apache-log4j-recommended-mitigation-measures are only acceptable until updates are available.

Weakness Enumeration

CWE-ID	CWE Name	Source
CWE-20	Improper Input Validation	 Apache Software Foundation
CWE-400	Uncontrolled Resource Consumption	 Apache Software Foundation
CWE-502	Deserialization of Untrusted Data	 Apache Software Foundation
CWE-917	Improper Neutralization of Special Elements used in an Expression Language Statement ('Expression La	 NIST

Known Affected Software Configurations

Configuration 1 ([hide](#))

 cpe:2.3:o:siemens:6bk1602-0aa12-0tp0_firmware:*:*:*:*:*:* Show matching CPE(s) ▼	Up to (excluding) 2.7.0
Running on/with	
cpe:2.3:h:siemens:6bk1602-0aa12-0tp0-*:*:*:*:*:* Show matching CPE(s) ▼	

Configuration 2 ([hide](#))


 cpe:2.3:o:siemens:6bk1602-0aa22-0tp0_firmware:*:*:*:*:*:* Show matching CPE(s) ▼	Up to (excluding) 2.7.0
Running on/with	
cpe:2.3:h:siemens:6bk1602-0aa22-0tp0-*:*:*:*:*:* Show matching CPE(s) ▼	

Configuration 3 ([hide](#))

 cpe:2.3:o:siemens:6bk1602-0aa32-0tp0_firmware:*:*:*:*:*:* Show matching CPE(s) ▼	Up to (excluding) 2.7.0
Running on/with	
cpe:2.3:h:siemens:6bk1602-0aa32-0tp0-*:*:*:*:*:* Show matching CPE(s) ▼	

Configuration 4 ([hide](#))

 cpe:2.3:o:siemens:6bk1602-0aa42-0tp0_firmware:*:*:*:*:*:* Show matching CPE(s) ▼	Up to (excluding) 2.7.0
Running on/with	



Early Results?

INFORMATIVE DEPLOYMENT LAUNCHED • AUGUST 25, 2026

4,461

**Records enriched
during the first week**


4.4%

Failure Rate


92.72K

**Average Tokens
per CVE**

What are we Automating?



What is the Solution Architecture?



How is Evidence Collected?

PUBLIC SOURCES

**Vendor
advisories**

Release notes

Patches

Issues


**Exploration
Agent**

Evidence package


**Reused by: References tagging,
CWE, CVSS, and Configurations**



How is Evidence Applied?



SPECIALIZED AGENTS



DETERMINISTIC CONTROLS



ANALYST REVIEW

AI proposes




**Software validates
and builds**



**Analysts decide when
uncertainty remains**

How is Traceable Output Produced?



AI
Software



Propose
Validate, Ground, Build

Traceable Output By Stage




SPECIALIST	INTERPRETS	BOUNDED BY	PRODUCES
Reference tagging	Source Purpose	Allowed tags	Tags + Evidence
CWE mapping	Weakness Mechanism	CWE catalogue	Weakness + Rationale
CVSS scoring	Exploitability + Impact	Policy + Calculation	Vector + Justification
Configuration	Product applicability	CPE + Configuration Rules	Scope + reasoning

Shared Context • Task-specific Instructions • Structured Outputs

How Many Workflows?


CVE Enrichment



CPE Generation




Operational




Why does CPE Need Its Own Workflow?

CVE Configuration



Logical applicability configuration


CPE Generation



Normalized CPE name and attributes

What does CPE Generation look like?

CPE Name Generation Workflow



What Did We Learn?

Lessons Learned:

- 01 **Web Evidence Collection**
- 02 **Ground decisions in evidence**
- 03 **Treat context as architecture**
- 04 **Decompose until work is bounded**
- 05 **Give QA the same evidence**
- 06 **Design integration as a workflow**

Why Is Web Evidence Collection Difficult?

LESSON 1 OF 6


OBSERVATION

Inconsistent across sources

IMPACT

Weak evidence means weak
CVE and CPE decisions

RESPONSE



Why Must Every Decision Cite Evidence?

LESSON 2 OF 6


OBSERVATION

Plausible proposal without accessible support

IMPACT

Difficult to trust, speculation of hallucination

RESPONSE



What Happens When Context is Lost?

LESSON 3 OF 6


OBSERVATION

Long context loses evidence and prior decisions

IMPACT

Lower accuracy, weaker explanation

RESPONSE



Why did Smaller Agent calls Perform Better?

LESSON 4 OF 6

OBSERVATION

One large call mixes evidence, rules, reasoning

IMPACT

Failure is difficult to diagnose

RESPONSE

Complex Stage

Precise call 1

Precise call 2

Can QA Succeed without Complete Context?

LESSON 5 OF 6


OBSERVATION

QA lacks evidence

IMPACT

More subjective interpretation

RESPONSE



Why Is Integration Difficult at Scale?

LESSON 6 OF 6


OBSERVATION

Queues, timeouts, schemas

IMPACT

False positive or incomplete proposals

RESPONSE



Additional Observations

Broader AI concerns that also affect a multi-agent workflow:

Model
Consistency

Model
Selection

Reasoning
Reliability

Prompt
Injection

What is next?

Evidence & feedback

Review calibration

Broader deployment



Improve Evidence collection

Investigate better retrieval mechanism

CPE Suggestions

Close the feedback loop

Establish Statistical Measures

Public Release

Questions?

