
GCVE-BCP-13 - Vulnerability Assigner Scorecard

GCVE.eu



Contents

1	GCVE-BCP-13 - Vulnerability Assigner Scorecard	1
1.1	Introduction	1
1.2	Goals	2
1.3	Design Principles	3
1.3.1	Score the Publisher's Own Contribution	3
1.3.2	Evaluate Usefulness, Not Format Compliance Alone	3
1.3.3	Use Graded Criteria	3
1.3.4	Do Not Require a Specific Severity Methodology	4
1.3.5	Keep Context Separate from the Score	4
1.3.6	Methodology Must Be Versioned	4
1.4	Terminology	5
1.4.1	Assigner	5
1.4.2	CNA	5
1.4.3	GNA	5
1.4.4	Record Score	5
1.4.5	Assigner Score	5
1.4.6	Pillar	6
1.4.7	Context Indicator	6
1.4.8	Ranked Assigner	6
1.4.9	Unranked Assigner	6
1.5	Requirements Language	6
1.6	Score Structure	6
1.7	Pillar 1 — Describe	7
1.7.1	Description — 10 points	7
1.7.2	Title — 5 points	8
1.7.3	CWE / Weakness — 10 points	8
1.8	Pillar 2 — Scope	9
1.8.1	Affected Product — 10 points	9
1.8.2	Affected Versions — 8 points	9
1.8.3	CPE / Machine-Readable Product Identification — 7 points	10

1.9	Pillar 3 — Assess	10
1.9.1	Severity — 15 points	11
1.9.2	Modern Severity — 5 points	11
1.9.3	Impact — 5 points	12
1.10	Pillar 4 — Act	12
1.10.1	References — 7 points	12
1.10.2	Reference Tags — 5 points	12
1.10.3	Remediation — 8 points	13
1.10.4	Solutions and Workarounds — 5 points	13
1.11	Grades	14
1.12	Evaluation Period	14
1.13	Ranking Threshold	15
1.14	Record Eligibility	15
1.15	CNA Evaluation	16
1.16	GNA Evaluation	16
1.17	External Enrichment	17
1.18	Context Indicators	17
1.18.1	External Severity Enrichment	18
1.18.2	External Weakness Enrichment	18
1.18.3	No Severity Statement	18
1.18.4	Short Descriptions	18
1.18.5	Known Exploited Vulnerabilities	18
1.18.6	Reservation-to-Publication Time	19
1.18.7	Rejected Records	19
1.18.8	Publication Volume	19
1.19	Timeliness	19
1.20	Comparative Scoreboards	20
1.21	Presentation and Improvement Signals	20
1.22	Badges	21
1.23	Rule-Level Transparency	21
1.24	Historical Trends	22
1.25	Methodology Versioning	23
1.26	BCP-13 Methodology Version 1	24
1.27	Reproducibility	24
1.28	Machine-Readable Output	25
1.29	Reference Implementation	25
1.30	Implementation Guidance	26
1.30.1	Producers	26

1.30.2 Scorecard Implementers	27
1.30.3 Consumers	27
1.31 Gaming and Metric Optimisation	27
1.32 Security Considerations	28
1.33 Fairness and Interpretation Considerations	29
1.34 Privacy Considerations	29
1.35 Interoperability	30
1.36 Future Evolution	30
1.37 Example	31
1.38 References	32
1.38.1 GCVE	32
1.38.2 Vulnerability-Lookup	32
1.38.3 Related Work	32
1.39 Acknowledgements	32

1 GCVE-BCP-13 - Vulnerability Assigner Scorecard



GCVE.eu

- **Version:** 1.0
- **Status:** Draft (for Public Review)
- **Date:** 2026-10-05
- **Authors:** GCVE Working Group
- **BCP ID:** BCP-13

This guide is distributed and available under [CC-BY-4.0](#).

Copyright (C) 2025-2026 GCVE Initiative.

1.1 Introduction

Vulnerability records are published by a large and diverse set of vulnerability assigners, including CVE Numbering Authorities (CNAs), GCVE Numbering Authorities (GNAs), vendors, security organisations, researchers and other vulnerability information producers.

Although these organisations may publish records using compatible formats, the completeness and operational usefulness of the resulting records can vary significantly.

This Best Current Practice defines a **Vulnerability Assigner Scorecard**: a transparent and reproducible methodology to evaluate how completely an assigner documents the vulnerability records it publishes.

The objective is not to judge the importance of an assigner, the quality of its vulnerability research, or the severity of the vulnerabilities it handles. Instead, the scorecard evaluates whether published records contain information useful to downstream consumers for four fundamental questions:

1. **Describe** — Can a reader understand what the vulnerability is?
2. **Scope** — Can a reader determine which products and versions are affected?
3. **Assess** — Can a reader understand the potential severity and impact?
4. **Act** — Can a reader determine what actions can be taken?

The scorecard is designed to support:

- vulnerability database operators;
- CNAs and GNAs wishing to evaluate their publication practices;
- vulnerability researchers;
- security teams consuming vulnerability data;
- automated vulnerability-management tooling;
- software vendors;
- vulnerability ecosystem researchers.

The methodology defined in this BCP is based on the operational scorecard implemented in **Vulnerability-Lookup**.

1.2 Goals

BCP-13 has the following goals:

- provide an objective and reproducible methodology for evaluating vulnerability record completeness;
- encourage improvement in the quality and usefulness of vulnerability advisories;
- make the scoring criteria understandable to vulnerability publishers;
- avoid rewarding information added by third parties as though it had been supplied by the original assigner;
- distinguish scoring from contextual observations;
- support both CNAs and GNAs;
- permit independent implementations;
- make changes to the methodology explicit and versioned;
- avoid penalising legitimate differences in vulnerability-management practices.

The scorecard is intended primarily as a tool for **continuous improvement and transparency**.

It MUST NOT be interpreted as a general measure of the competence, trustworthiness or authority of an organisation.

1.3 Design Principles

1.3.1 Score the Publisher's Own Contribution

The scorecard SHOULD evaluate information authored or published by the assigner being evaluated.

Information subsequently added by another organisation, enrichment service, Authorized Data Publisher (ADP), vulnerability database or automated process MUST NOT be credited to the original assigner.

For CVE JSON 5 records, the score SHOULD therefore be calculated from the CNA container when evaluating a CNA.

For a GNA publishing its own records, the content attributable to that GNA SHOULD be evaluated.

External enrichment MAY be reported separately as contextual information.

1.3.2 Evaluate Usefulness, Not Format Compliance Alone

A record can be syntactically valid while providing very little useful information.

BCP-13 therefore evaluates the presence and quality of information rather than simple schema compliance.

For example, a non-empty one-line description SHOULD NOT necessarily receive the same credit as a sufficiently descriptive vulnerability explanation.

1.3.3 Use Graded Criteria

Rules SHOULD support partial credit where useful.

A scoring system based exclusively on binary presence or absence can create misleading incentives. For example:

- a very short description and a detailed description are not equivalent;
- identifying only a product but not its vendor is less useful than identifying both;
- supplying a mitigation reference is useful even when a direct patch reference is unavailable.

1.3.4 Do Not Require a Specific Severity Methodology

Different vulnerability publishers use different severity-assessment methods.

BCP-13 therefore considers a meaningful severity statement valid regardless of whether it is expressed using:

- CVSS;
- SSVC;
- a vendor severity classification;
- another documented severity methodology.

Modern machine-readable severity information MAY receive additional credit, but the scorecard MUST NOT require CVSS as the only acceptable severity methodology.

1.3.5 Keep Context Separate from the Score

Some information is valuable for understanding an assigner's publication practice but cannot be fairly converted into a score.

Such information SHOULD be displayed as contextual indicators rather than included in the numerical score.

Examples include:

- publication delay;
- external enrichment;
- rejected records;
- known exploited vulnerabilities;
- publication volume.

1.3.6 Methodology Must Be Versioned

The scoring methodology MUST be versioned.

Any change affecting:

- a rule;
- a rule weight;

- partial-credit criteria;
- grade boundaries;
- ranking thresholds;
- default evaluation windows;

MUST produce a new methodology version.

A published score SHOULD always identify the methodology version under which it was calculated.

1.4 Terminology

1.4.1 Assigner

An organisation or entity responsible for assigning or publishing vulnerability identifiers and their associated vulnerability records.

Examples include a CNA or a GNA.

1.4.2 CNA

A **CVE Numbering Authority** participating in the CVE Program.

1.4.3 GNA

A **GCVE Numbering Authority** participating in the GCVE ecosystem.

1.4.4 Record Score

The numerical score calculated for one vulnerability record according to the methodology defined in this BCP.

1.4.5 Assigner Score

The arithmetic mean of the record scores attributable to an assigner for the selected evaluation period.

1.4.6 Pillar

A logical group of scoring rules representing one aspect of vulnerability record usefulness.

BCP-13 defines four pillars:

- Describe;
- Scope;
- Assess;
- Act.

1.4.7 Context Indicator

A metric displayed alongside the score which does not contribute points to the score.

1.4.8 Ranked Assigner

An assigner with sufficient publication volume during the selected evaluation window to participate in comparative ranking.

1.4.9 Unranked Assigner

An assigner for which a score can be calculated but which has not published enough records during the evaluation period to satisfy the ranking threshold.

1.5 Requirements Language

The key words **MUST**, **MUST NOT**, **REQUIRED**, **SHOULD**, **SHOULD NOT**, **MAY**, and **OPTIONAL** in this document are to be interpreted as described in RFC 2119 and RFC 8174 when, and only when, they appear in all capitals.

1.6 Score Structure

Each vulnerability record is evaluated using four pillars:

Pillar	Maximum Points
Describe	25
Scope	25
Assess	25
Act	25
Total	100

Each rule returns a credit between 0.0 and 1.0.

The points awarded for a rule are:

```
1 awarded_points = rule_weight × rule_credit
```

The record score is:

```
1 record_score = sum(awarded_points)
```

The resulting score ranges from 0 to 100.

The assigner score for an evaluation period is the arithmetic mean of all eligible record scores:

```
1 assigner_score =  
2   sum(record_score for eligible records)  
3   / number_of_eligible_records
```

An implementation SHOULD retain sufficient precision internally and MAY round values for presentation.

1.7 Pillar 1 — Describe

The **Describe** pillar evaluates whether a vulnerability record provides enough information to understand what the vulnerability is.

Maximum: **25 points**.

1.7.1 Description — 10 points

The record SHOULD contain an English-language description which meaningfully describes the vulnerability.

The following graded credit is RECOMMENDED:

Description Length	Credit
200 characters or more	1.0
80–199 characters	0.7
40–79 characters	0.4
less than 40 characters	0.0

Descriptions identified as placeholders SHOULD receive zero credit.

Examples of placeholder-like values include information equivalent to:

- 1 TODO
- 2 TBD
- 3 N/A
- 4 Reserved
- 5 No description available

Implementations MAY maintain a documented set of recognised placeholders.

1.7.2 Title — 5 points

A meaningful vulnerability title earns full credit.

Condition	Credit
Meaningful title present	1.0
No title	0.0

A title SHOULD provide a concise identification of the vulnerability rather than simply repeat its identifier.

1.7.3 CWE / Weakness — 10 points

A weakness expressed using a CWE identifier SHOULD receive full credit.

Condition	Credit
CWE identifier present	1.0
Free-text weakness or problem type without CWE	0.25
No weakness information	0.0

The use of CWE is encouraged because it provides a machine-readable and interoperable weakness taxonomy.

1.8 Pillar 2 — Scope

The **Scope** pillar evaluates whether consumers can determine which products and versions are affected.

Maximum: **25 points**.

1.8.1 Affected Product — 10 points

An affected entry SHOULD identify both the vendor and product.

Condition	Credit
Vendor and product identified	1.0
Only vendor or product identified, or one is n/a	0.5
Neither meaningfully identified	0.0

1.8.2 Affected Versions — 8 points

The record SHOULD provide information permitting consumers to determine affected versions.

Full credit SHOULD be given when an affected entry contains:

- an explicit affected version;
- an affected version range;

- lower or upper version bounds;
- another machine-readable equivalent.

A `defaultStatus` without explicit version information MAY receive partial credit.

Condition	Credit
Explicit version or version range	1.0
<code>defaultStatus</code> only	0.5
No useful version information	0.0

1.8.3 CPE / Machine-Readable Product Identification — 7 points

A machine-readable product identifier SHOULD be supplied where possible.

Full credit is awarded when the record includes:

- a CPE associated with an affected product; or
- a `cpeApplicability` structure or equivalent machine-readable applicability expression.

Condition	Credit
Machine-readable CPE/applicability present	1.0
Not present	0.0

Future methodology versions MAY support additional machine-readable product identification schemes such as PURL or equivalent identifiers.

Such changes MUST result in a new methodology version if they affect scoring.

1.9 Pillar 3 — Assess

The **Assess** pillar evaluates whether a consumer can understand how serious the vulnerability may be and what impact exploitation can have.

Maximum: **25 points**.

1.9.1 Severity — 15 points

The record SHOULD contain a severity statement authored by the assigner.

Full credit SHOULD be awarded for a recognised severity assessment such as:

- CVSS v3.x;
- CVSS v4.x;
- SSVC;
- a vendor-defined severity expressed as an **other** metric or equivalent structured field.

CVSS v2 alone SHOULD receive partial credit.

Condition	Credit
CVSS v3/v4, SSVC, vendor severity or equivalent modern assessment	1.0
CVSS v2 only	0.5
No severity statement	0.0

The absence of CVSS MUST NOT by itself imply that a vulnerability record lacks a severity assessment.

1.9.2 Modern Severity — 5 points

Additional credit SHOULD be awarded for a modern structured severity or decision methodology.

Current examples include:

- CVSS v4;
- SSVC.

Condition	Credit
Modern severity or decision methodology present	1.0
Not present	0.0

This rule deliberately distinguishes between providing **some meaningful severity statement** and adopting newer structured severity methodologies.

1.9.3 Impact — 5 points

The record **SHOULD** describe the consequence or impact of successful exploitation.

Full credit **SHOULD** be awarded when the record contains:

- a CAPEC identifier; or
- a meaningful impact description in the appropriate structured field.

Condition	Credit
Impact or CAPEC information present	1.0
Not present	0.0

1.10 Pillar 4 — Act

The **Act** pillar evaluates whether the record provides enough information for a consumer to investigate, mitigate or remediate the vulnerability.

Maximum: **25 points**.

1.10.1 References — 7 points

At least one relevant external reference **SHOULD** be supplied.

Condition	Credit
One or more references	1.0
No references	0.0

1.10.2 Reference Tags — 5 points

References **SHOULD** be semantically classified where supported by the record format.

Condition	Credit
At least one tagged reference	1.0
References exist but none are tagged	0.0

Tags improve automated processing by distinguishing, for example:

- vendor advisories;
- patches;
- release notes;
- technical reports;
- proof-of-concept material;
- issue trackers.

1.10.3 Remediation — 8 points

A record SHOULD contain a reference to actionable remediation information.

The following credit is RECOMMENDED:

Condition	Credit
Reference tagged <code>patch</code>	1.0
Reference tagged <code>vendor-advisory</code> , <code>mitigation</code> , or <code>release-notes</code>	0.6
No remediation reference	0.0

1.10.4 Solutions and Workarounds — 5 points

A record SHOULD contain a solution, mitigation or workaround where one is known.

Condition	Credit
Solution or workaround present	1.0

Condition	Credit
Not present	0.0

The absence of a solution MUST NOT be interpreted as evidence that the assigner failed to coordinate remediation. A solution may legitimately be unavailable at the time of publication.

The rule measures the contents of the published record, not the underlying disclosure process.

1.11 Grades

For human-readable presentation, implementations MAY translate the numerical score into a grade.

The following boundaries define the BCP-13 version 1 grade profile:

Score	Grade
90–100	A
75–<90	B
60–<75	C
45–<60	D
<45	E

Grades are presentation aids.

The numerical score and methodology version SHOULD always remain available so that consumers are not required to rely exclusively on the grade.

1.12 Evaluation Period

A score SHOULD be associated with a clearly defined evaluation period.

For comparative evaluation of current publication practices, BCP-13 RECOMMENDS a rolling window of:

1 180 days

The rolling window helps ensure that the score reflects current practice rather than publication practices from many years earlier.

Implementations MAY additionally calculate:

- all-time scores;
- scores per publication year;
- scores for arbitrary periods.

Every published score MUST state the evaluation period used.

1.13 Ranking Threshold

A numerical score can be calculated from very few records, but such a score may not meaningfully represent an assigner's normal publication practice.

An assigner SHOULD therefore publish at least:

1 10 records

within the evaluation window before being included in a comparative ranking.

An assigner below this threshold SHOULD:

- still receive a score;
- remain visible;
- be identified as **unranked**.

This distinction prevents very low publication volume from creating misleading comparative results.

Both the evaluation-window duration and minimum-record threshold MUST be included in machine-readable scorecard results.

1.14 Record Eligibility

Only records attributable to the assigner being evaluated SHOULD contribute to its score.

Rejected records SHOULD NOT normally be included when computing the completeness score of published vulnerability advisories.

Rejected records MAY be reported separately as contextual information.

An implementation MUST document:

- which records are eligible;
- which publication timestamp is used;
- how modified records are treated;
- how rejected or withdrawn records are handled.

1.15 CNA Evaluation

When evaluating a CNA using CVE JSON 5 data, scoring MUST be based on information supplied by that CNA.

Information added later by another container or data provider MUST NOT increase the CNA's score.

For example, if:

1. a CNA publishes a record without CWE information; and
2. an ADP subsequently adds a CWE;

the CNA MUST NOT receive CWE points for that record.

The external CWE MAY instead contribute to an **external enrichment** context indicator.

This separation is important because the scorecard measures the publication practice of the assigner rather than the final aggregate quality of a vulnerability record after enrichment.

1.16 GNA Evaluation

GNAs SHOULD be evaluated according to the same fundamental scoring methodology as CNAs when they publish compatible vulnerability records.

The scorecard SHOULD evaluate information attributable to the GNA itself.

A GNA-specific implementation MAY additionally expose contextual information including:

- the associated GCVE identifier;
- corresponding CVE identifiers;
- relationships to records published by other GNAs;

- identifier placement or format-version information;
- use of GCVE extensions.

GNA-specific contextual indicators MUST NOT alter the core BCP-13 score unless explicitly introduced by a future methodology version.

1.17 External Enrichment

Vulnerability information is frequently enriched after initial publication.

Examples include:

- CWE classification;
- CVSS calculation;
- CPE mapping;
- exploitation status;
- KEV status;
- references;
- analyst annotations.

Such enrichment is valuable but SHOULD remain distinguishable from the original assigner's work.

An implementation SHOULD therefore expose enrichment metrics separately.

For example:

1	CNA severity present:	64%
2	Severity supplied externally:	31%
3	No severity available:	5%

The distinction allows consumers to understand both:

1. the completeness of the original publisher's records; and
2. the extent to which downstream consumers depend on third-party enrichment.

1.18 Context Indicators

The following contextual information is RECOMMENDED but MUST NOT contribute directly to the BCP-13 score.

1.18.1 External Severity Enrichment

Percentage of records for which:

- the assigner supplied no severity; and
- a third party subsequently supplied one.

1.18.2 External Weakness Enrichment

Percentage of records for which:

- the assigner supplied no CWE; and
- a third party subsequently supplied one.

1.18.3 No Severity Statement

Percentage of records containing no severity statement from any accepted method.

1.18.4 Short Descriptions

Percentage of records with descriptions below a defined length.

The current reference implementation uses:

1	100	characters
---	-----	------------

as a contextual threshold.

1.18.5 Known Exploited Vulnerabilities

An implementation MAY report:

- how many vulnerabilities published by the assigner appear in one or more Known Exploited Vulnerability catalogues;
- how many of those records were published without a severity statement.

Known exploitation MUST NOT automatically modify the assigner's BCP-13 score.

1.18.6 Reservation-to-Publication Time

An implementation MAY display:

- median reservation-to-publication delay;
- 90th percentile reservation-to-publication delay.

This metric **MUST NOT** contribute to the score.

A long period between identifier reservation and publication may result from legitimate coordinated vulnerability disclosure, embargo requirements, vendor remediation schedules or other circumstances that cannot be inferred from timestamps alone.

1.18.7 Rejected Records

The number or proportion of rejected records MAY be displayed as context.

1.18.8 Publication Volume

The number of eligible records **SHOULD** always be displayed alongside an assigner score.

1.19 Timeliness

BCP-13 deliberately does **not** assign points for publication speed.

The time between identifier reservation and publication can represent many different processes, including:

- responsible disclosure;
- an embargo;
- patch development;
- coordination between multiple vendors;
- supply-chain coordination;
- regulatory requirements;
- delayed publication;
- administrative backlog.

Without additional information, these situations cannot reliably be distinguished.

Timeliness SHOULD therefore be observable but MUST NOT be interpreted automatically as a quality score.

1.20 Comparative Scoreboards

An implementation MAY publish a scoreboard comparing multiple assigners.

A scoreboard SHOULD provide, at minimum:

- assigner identifier or name;
- numerical score;
- grade, if grades are used;
- number of evaluated records;
- evaluation period;
- methodology version;
- ranked or unranked state.

Assigners satisfying the minimum publication threshold MAY be ordered by numerical score.

Unranked assigners SHOULD remain identifiable as such and SHOULD NOT be presented as directly comparable to ranked assigners.

1.21 Presentation and Improvement Signals

Implementations MAY provide additional views identifying particularly complete publication practices or recurrent documentation gaps.

Such views are presentation features and are not part of the core scoring methodology.

For example, an implementation MAY identify objective conditions such as:

- all eligible records lacking a severity statement;
- all eligible records lacking CWE information;
- external severity enrichment required for more than half of records;
- descriptions below a defined length for more than 30% of records;

- missing affected product information for more than half of records;
- known exploited vulnerabilities published without severity information.

Any such signal:

- MUST be based on explicit and documented criteria;
- MUST identify the evaluation window;
- MUST distinguish facts derived from records from qualitative interpretation;
- SHOULD provide enough information for the publisher to identify what can be improved.

Implementations SHOULD favour terminology oriented toward improving vulnerability data quality rather than implying that the score represents the overall quality or competence of an organisation.

1.22 Badges

Implementations MAY provide machine-generated badges representing scorecard results.

A badge SHOULD include at least:

- the grade or numerical score;
- an indication when the assigner is unranked.

The badge SHOULD link to a page providing:

- the complete score;
- methodology version;
- evaluation period;
- evaluated-record count;
- rule-level results.

A badge MUST NOT be treated as a cryptographic or trust assertion.

1.23 Rule-Level Transparency

Implementations SHOULD expose more than the aggregate score.

For every assigner, it SHOULD be possible to inspect:

- credit rate per rule;
- points awarded per rule;
- score per pillar;
- overall score;
- record count;
- context indicators.

This allows the scorecard to function as an improvement tool rather than merely a ranking mechanism.

For example:

```
1 Describe: 21.8 / 25
2 Scope:    17.2 / 25
3 Assess:   14.9 / 25
4 Act:      22.6 / 25
5
6 Overall:  76.5 / 100
7 Grade:    B
```

A publisher can then determine which aspects of its records could be improved.

1.24 Historical Trends

Implementations SHOULD retain sufficient information to show how an assigner's publication practice evolves.

Useful historical views include:

- score by calendar year;
- score by methodology version;
- rule-level success rates over time;
- pillar scores over time.

Scores calculated under different methodology versions SHOULD NOT be compared without clearly identifying the change in methodology.

An implementation MUST NOT silently reinterpret historical scores using newer rules while presenting them as though they were originally calculated under those rules.

1.25 Methodology Versioning

Each scorecard calculation MUST identify a methodology version.

A methodology version MUST change whenever there is a scoring-significant change, including:

- addition or removal of a rule;
- modification of rule credit;
- modification of a rule weight;
- modification of grade boundaries;
- modification of ranking eligibility criteria;
- modification of the default evaluation period.

A result SHOULD contain information equivalent to:

```
1 {  
2   "methodology": 1,  
3   "window_days": 180,  
4   "minimum_records": 10  
5 }
```

An implementation MAY additionally expose:

- software name;
- software version;
- calculation timestamp;
- dataset version;
- source identifier.

Example:

```
1 {  
2   "methodology": 1,  
3   "implementation": "Vulnerability-Lookup",  
4   "implementation_version": "X.Y.Z",  
5   "calculated_at": "2026-09-30T12:00:00Z",  
6   "window_days": 180,  
7   "minimum_records": 10  
8 }
```

1.26 BCP-13 Methodology Version 1

The initial BCP-13 methodology is defined as:

- 13 graded rules;
- 4 pillars;
- 25 points per pillar;
- maximum score of 100;
- assigner-authored information only;
- any meaningful severity statement accepted;
- CVSS v4 or SSVC receives modern-severity credit;
- timeliness displayed as context rather than scored;
- default ranking window of 180 days;
- minimum of 10 publications for ranking.

The methodology identifier is:

1 BCP-13 methodology version 1

1.27 Reproducibility

A BCP-13 implementation SHOULD make its score reproducible by another implementation operating on the same input dataset.

Implementations SHOULD therefore document:

- record selection;
- publication-window calculation;
- treatment of updates;
- field extraction;
- placeholder detection;
- supported severity formats;
- supported CWE representation;
- affected-product interpretation;

- version-range interpretation;
- remediation tags;
- rounding behaviour.

Where an implementation introduces behaviour not defined by the current BCP methodology, that behaviour **MUST NOT** silently change the normative BCP-13 score.

1.28 Machine-Readable Output

Implementations **SHOULD** provide a machine-readable representation of scorecard results.

A scorecard result **SHOULD** contain information equivalent to:

```
1 {
2   "assigner": "example",
3   "kind": "cna",
4   "methodology": 1,
5   "period": {
6     "type": "rolling-window",
7     "days": 180
8   },
9   "minimum_records": 10,
10  "record_count": 42,
11  "ranked": true,
12  "score": 82.4,
13  "grade": "B",
14  "pillars": {
15    "describe": 22.1,
16    "scope": 19.7,
17    "assess": 18.4,
18    "act": 22.2
19  }
20 }
```

Implementations **MAY** expose additional rule-level and contextual fields.

1.29 Reference Implementation

The initial reference implementation of BCP-13 is available in **Vulnerability-Lookup**.

The human-readable scorecard is available at:

```
1 https://vulnerability.circl.lu/assigners/scorecard
```

GNA scores can be selected from the same interface.

The methodology documentation is available at:

```
1 https://www.vulnerability-lookup.org/documentation/scorecard.html
```

The reference implementation also provides API endpoints for programmatic access.

At the time of publication of this BCP, these include:

```
1 GET /api/stats/scorecard/ranking
2 GET /api/stats/scorecard/gnas
3 GET /api/stats/scorecard/assigner
4 GET /api/stats/scorecard/signals
```

BCP-13 does not require implementations to use these exact HTTP paths.

They document the API of the reference implementation rather than the protocol defined by this BCP.

1.30 Implementation Guidance

1.30.1 Producers

Vulnerability publishers wishing to improve their BCP-13 score SHOULD focus primarily on publishing useful information rather than attempting to optimise specifically for the score.

In particular, publishers SHOULD provide where possible:

- meaningful titles and descriptions;
- CWE classifications;
- explicit vendor and product information;
- affected versions or version ranges;
- machine-readable product identifiers;
- a severity assessment;
- impact information;
- useful references with semantic tags;
- remediation references;
- solutions or workarounds.

The scorecard SHOULD act as a reflection of good vulnerability publication practice rather than become an objective independent of that practice.

1.30.2 Scorecard Implementers

Implementers SHOULD:

- use deterministic rules;
- publish the methodology;
- expose methodology versions;
- preserve the distinction between publisher information and external enrichment;
- show record counts;
- distinguish ranked and unranked assigners;
- expose contextual indicators separately from score components;
- permit assigners to reproduce or inspect their results.

1.30.3 Consumers

Consumers SHOULD NOT use a BCP-13 score as the sole criterion for determining:

- whether vulnerability information should be trusted;
- whether an organisation is a competent CNA or GNA;
- whether a vulnerability is valid;
- whether a vulnerability is important;
- whether an assigner should participate in a vulnerability ecosystem.

The score reflects the **completeness of published vulnerability records according to a documented methodology**.

It does not measure every aspect of vulnerability coordination or publication quality.

1.31 Gaming and Metric Optimisation

Any public metric can influence behaviour.

Implementations and consumers SHOULD therefore distinguish between satisfying a scoring condition and providing genuinely useful vulnerability information.

Examples of undesirable metric optimisation include:

- adding meaningless text solely to exceed description-length thresholds;
- assigning arbitrary CWE identifiers;
- publishing incorrect CPEs solely to obtain CPE credit;
- publishing unjustified severity ratings;
- tagging unrelated references as patches;
- inserting non-actionable workaround text.

A higher score obtained through inaccurate information is contrary to the objectives of BCP-13.

Correctness SHOULD take precedence over completeness.

When reliable information is unavailable, omitting a field can be preferable to publishing unsupported information.

1.32 Security Considerations

Scorecard systems process data from potentially untrusted vulnerability records.

Implementations MUST treat record content as untrusted input.

In particular, implementations SHOULD protect against:

- HTML or script injection;
- malicious URLs;
- excessively large fields;
- malformed structured data;
- pathological inputs intended to consume excessive processing resources;
- manipulated timestamps;
- identifier spoofing.

Scorecard computation SHOULD operate on validated vulnerability records wherever possible.

1.33 Fairness and Interpretation Considerations

Different assigners operate under different constraints.

Some publish:

- vendor advisories;
- open-source project vulnerabilities;
- third-party vulnerability reports;
- vulnerability records reconstructed from historical archives;
- vulnerabilities affecting products with limited version metadata.

Scorecard results SHOULD therefore be interpreted in context.

A lower score does not necessarily imply that an assigner performed poor vulnerability coordination, and a higher score does not guarantee that every field is correct.

The scorecard evaluates observable publication characteristics.

It cannot directly evaluate:

- accuracy of undisclosed information;
- quality of researcher communication;
- quality of vendor coordination;
- embargo handling;
- remediation engineering;
- exploitability;
- vulnerability validity;
- organisational competence.

1.34 Privacy Considerations

BCP-13 scores organisations or vulnerability-information producers rather than individual researchers.

Implementations SHOULD avoid unnecessarily exposing personal information.

Where an assigner is an individual, implementations SHOULD limit displayed information to information already required for attribution or publication.

1.35 Interoperability

BCP-13 is designed to be largely independent of vulnerability identifier namespace.

It may be applied to:

- CVE records;
- GCVE records;
- other vulnerability formats;

provided that the implementation can reliably map the required semantic fields.

An implementation using a non-CVE representation **MUST** document how its fields map to the BCP-13 rules.

The methodology **SHOULD** remain independent from the serialization format whenever equivalent semantic information is available.

1.36 Future Evolution

Future BCP-13 methodology versions may consider additional structured information such as:

- PURL identifiers;
- SBOM applicability;
- VEX information;
- EPSS or exploitation-related data;
- disclosure timelines;
- provenance;
- credit information;
- richer remediation semantics;
- supply-chain relationships;
- GCVE extension usage.

New information **SHOULD** only influence the numerical score when:

1. it has clear operational value;
2. it can be evaluated reproducibly;

3. it does not unfairly favour one vulnerability-publication model;
4. the scoring change is introduced through a new methodology version.

Contextual indicators are preferred when these conditions cannot be met.

1.37 Example

Consider an assigner publishing 50 vulnerability records during the last 180 days.

Its aggregated pillar results might be:

Pillar	Score
Describe	23.1 / 25
Scope	18.7 / 25
Assess	20.4 / 25
Act	21.8 / 25
Overall	84.0 / 100

Using methodology version 1:

```
1 Score: 84.0
2 Grade: B
3 Records: 50
4 Window: 180 days
5 Ranking threshold: 10
6 Status: Ranked
7 Methodology: BCP-13 v1
```

Context information could additionally indicate:

```
1 Severity supplied by external enrichment: 12%
2 Weakness supplied by external enrichment: 8%
3 No severity statement: 4%
4 Descriptions below 100 characters: 6%
5 Median reservation-to-publication time: 38 days
6 90th percentile reservation-to-publication time: 121 days
7 Rejected records: 3
```

These contextual values do not modify the numerical score.

1.38 References

1.38.1 GCVE

- GCVE Best Current Practices: <https://gcve.eu/bcp/>
- GCVE-BCP-05 — GCVE Vulnerability Format: <https://gcve.eu/bcp/gcve-bcp-05/>
- GCVE-BCP-06 — Requirements and Evaluation Criteria for GCVE Numbering Authorities: <https://gcve.eu/bcp/gcve-bcp-06/>

1.38.2 Vulnerability-Lookup

- Vulnerability-Lookup Assigner Scorecard documentation: <https://www.vulnerability-lookup.org/documentation/scorecard.html>
- [Vulnerability-Lookup operational scorecard](#)
- Vulnerability-Lookup project: <https://github.com/cve-search/vulnerability-lookup>

1.38.3 Related Work

The concept of evaluating vulnerability assigners based on record completeness predates BCP-13.

The Vulnerability-Lookup implementation was inspired in part by the [CNA ScoreCard project](#).

BCP-13 differs by defining a graded methodology, separating assigner-authored data from third-party enrichment, supporting GNAs, accepting multiple severity methodologies, introducing a publication-volume threshold and requiring explicit methodology versioning.

1.39 Acknowledgements

BCP-13 builds on operational experience from the Vulnerability-Lookup project and on feedback from vulnerability publishers, CNAs, GNAs, CSIRTs, vendors, researchers and vulnerability-management practitioners.

Contributions, implementation feedback and proposals for future methodology revisions are welcome through the GCVE public development process.