

The Criterica Score Specification

The Criterica Score is a number from 0 to 100 attached to one legal asset, portfolio or counterparty at a stated date, with its target, horizon and reference population stated beside it. This is the specification, issued for comment.

[COMMENT ON THE SPECIFICATION](#)

[SEE THE PREREGISTRATION REGISTRY](#)

SERIES

Standard STD-002

VERSION

1.0-rc

PUBLISHED

9 October 2026

STATUS

Request for comment

AUTHORS

Criterica Intelligence Research

Corresponding author: Jack Cosentino

01 Purpose and scope

The Criterica Score is a comparable, versioned measure attached to a file. It exists so that two parties looking at the same legal asset, on the same date, against the same stated population, read the same number and know exactly what the number means.

This specification defines what the Score is, what it estimates, what it must state to be valid, when it must not be issued and how it changes over time. It does not publish the models that produce it.

[COMMENT ON THE SPECIFICATION](#)

- Funders and lenders pricing, financing or marking a legal asset.
 - Law firms assessing their own portfolio or a matter within it.
 - Insurers and reinsurers reserving or pricing claims exposure.
 - Platforms and auditors who need a number that can be reproduced and challenged.
-

02 Definition

A Criterica Score is a number from 0 to 100 attached to one legal asset, one portfolio or one counterparty at a stated date.

A valid Score always travels with five statements: the asset it describes, the date it is attached to, the outcome it targets, its horizon and its reference population. A number that arrives without them is not a Criterica Score.

1. Subject: the legal asset, portfolio or counterparty.
2. As-of date: the date the file was read.
3. Target: the outcome the Score estimates, as defined in the next section.
4. Horizon: the period the Score covers.
5. Reference population: the population the file is compared with.

The same file, scored again under the same version, on the same as-of date, returns the same number.

03 Target

The Score estimates the probability-weighted standing of the asset against its reference population on a stated outcome. The outcome is one of four questions: who prevails, how long it takes, what it

COMMENT ON THE SPECIFICATION

Standing is a relative position, not a forecast of one result. A high Score means the file sits high in its reference population on the stated outcome. It does not say that a given matter will win, resolve by a given date or collect a given amount.

The target is stated per asset class in the asset-class annex that accompanies each version of this specification. An asset class with no annex has no Score.

Outcome families

QUESTION	WHAT STANDING IS MEASURED ON
Who prevails	The probability-weighted standing on the stated result for the party the asset depends on.
How long it takes	The probability-weighted standing on time to the stated resolution event.
What it costs	The probability-weighted standing on the stated cost measure.
What it is worth	The probability-weighted standing on the stated value measure, shown as a labeled distribution.

04 Horizon

The horizon is the period the Score covers, counted from the as-of date. A Score always states its horizon, and a Score is only comparable with another Score that states the same horizon.

A horizon the models cannot support is not rounded to the nearest one they can. The file is returned as Not rated for that horizon.

05 Reference population

COMMENT ON THE SPECIFICATION

type and period.

A Score without a stated reference population is invalid. Changing the reference population changes the Score, so it is recorded with the number every time.

06 Mapping and grades

Model outputs map to the 0 to 100 scale by a monotone function: a file that stands higher on the stated outcome never receives a lower Score than a file that stands lower, all else equal. The mapping is calibrated against the reference population so that the scale means the same thing from one scoring run to the next.

The mapping is fixed within a version. It is re-mapped only at a version change, and the change is recorded in the change log.

Grades

Grades are five named bands on the scale. A grade describes position in the reference population and nothing else. It makes no claim about how well the Score performs.

BAND	NAME	SCALE	READING
Band 1	Well below reference	0 to 19	The file stands well below the reference population on the stated outcome.
Band 2	Below reference	20 to 39	The file stands below the reference population.
Band 3	Near reference	40 to 59	The file stands close to the middle of the reference population.
Band 4	Above reference	60 to 79	The file stands above the reference population.
Band 5	Well above reference	80 to 100	The file stands well above the reference population.

07 Not rated

When support is insufficient, the file is Not rated. It is never given a low Score. A low Score says the file stands low; Not rated says the Score cannot be issued.

A file is Not rated when any of these holds:

- The data on the file is too thin to support the stated outcome.
- The file falls outside the stated reference population.
- The sealed-outcome rule is violated.
- The stated horizon is not supported.

A Not rated result states which condition applied. It is recorded in the same way as a Score.

COMMENT ON THE SPECIFICATION

The inputs to a Score are the facts of the file as they stood at the as-of date: the parties, the forum, the claim or case type, the procedural history, the terms of the asset and the evidence supplied for it.

The sealed-outcome rule: the file scored carries no resolution field, no collected amount, no resolution date and no field that encodes them. A file that carries any of these is Not rated. The rule keeps the answer out of the question.

Inputs the scorer supplies and inputs the scorer assumes are labeled separately on the output. A seller-stated figure is tagged as seller-stated.

09 Validation status

Validation is reported per asset class. A claim about how well the Score performs for an asset class activates only after a preregistered backtest on that class passes, and the claim is made together with its evidence. Preregistrations and their results are published, whether they pass or fail, in the preregistration registry.

ASSET CLASS	STATUS
Pre-settlement consumer	Not yet validated; claims activate after a preregistered backtest passes
Law-firm portfolio	Not yet validated; claims activate after a preregistered backtest passes
Commercial litigation	Not yet validated; claims activate after a preregistered backtest passes
Medical receivables	Not yet validated; claims activate after a preregistered backtest passes
Insurance claims	Not yet validated; claims activate after a preregistered backtest passes

No row in this table says validated. A row changes only when a preregistered backtest for that asset class has passed and its result is posted.

[Read the preregistration registry.](#)

COMMENT ON THE SPECIFICATION

10 Versioning

The specification and the Score it defines carry semantic versions: a major version when the meaning of the Score changes, a minor version when a method is added or tightened without changing meaning, and a patch version for corrections to wording. A release candidate carries the suffix rc until comments are answered.

A change log records every version, its date and what changed. Each of these forces a re-mapping and therefore a new version:

- A change to the target or the outcome definition.
- A change to the reference population criteria.
- A change to the mapping function or to the band edges.
- A retrained model that moves the calibration.

Prior scores are not rewritten. A Score issued under a version keeps that version and stays valid as a record of what was issued. A file re-scored under a new version is issued as a new Score with its own version, and both appear on the record.

Change log

VERSION	DATE	CHANGE
1.0-rc	2026-10-09	Issued as a request for comment.

11 What the Score is not

- It is not a merits opinion on any matter.
- It is not legal advice.
- It is not a forecast of realized returns or of a case's realized dollar value.

COMMENT ON THE SPECIFICATION

Dollar outputs that accompany a Score are labeled distributions, and inputs are tagged by source. Where a return such as IRR or MOIC is shown, it is computed from those distributions and the stated terms of the asset, and it is labeled as computed.

12 Governance and comment

Changes to the specification are proposed in writing, discussed against the evidence, issued first as a release candidate and recorded in the change log. Anyone may comment on a release candidate through the form on this page. Each comment is read, answered in writing and, where it changes the text, credited in the change log by role and organization type.

The comment period for this release candidate is open. Its close date will be stated here when it is set.

Errata

Errata correct errors in a published version without changing its meaning. They are listed here with the date and the text corrected.

No errata have been issued for version 1.0-rc.

REQUEST FOR COMMENT

Comment on the specification.

Say what is unclear, what is missing or what you would change, and why. Each comment is answered in writing.

COMMENT ON THE SPECIFICATION

Criteria Intelligence Research (2026). The Criteria Score Specification. Standard STD-002, version 1.0-rc. Criteria. <https://critericaintelligence.com/score>

```
@techreport{criterica_std002_2026,
  author = {{Criteria Intelligence Research} and Cosentino, Jack},
  title = {The Criteria Score Specification},
  institution = {Criteria},
  type = {Standard},
  number = {STD-002},
  year = {2026},
  month = oct,
  note = {Version 1.0-rc},
  url = {https://critericaintelligence.com/score}
}
```

DOI: pending

See the Score applied to your file.

REQUEST AN AUDIT

PREREGISTRATION REGISTRY

FAQ

Questions



KEEP READING

Preregistration registry Methodology Model registry Research Duration intelligence

COMMENT ON THE SPECIFICATION