Also staged: Review No. 1: Microsoft 365 E5, page preview (live on the site from 12 October)

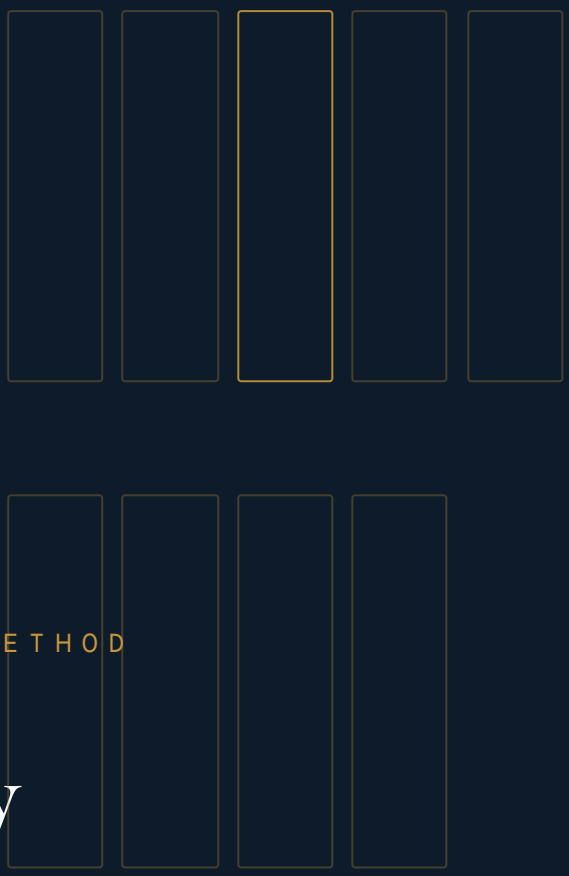

ALTLIBRE · SCORECARD METHOD

The Sovereignty Scorecard

A public method for reviewing the technology Europe depends on. Eight tests of jurisdiction, control and sustainability, published so that every score can be argued with.

ADRIAN HOLLISTER · ALTLIBRE · VERSION 1.2, 1 OCTOBER 2026

IN BRIEF

What this method is

Large analyst houses score what a vendor can execute. Their customers are largely the vendors. This page publishes the method AltLibre will use to score something they cannot: what the vendor owes the European citizen, and whether Europe can afford the relationship for the decade ahead.

Eight criteria, tested against checkable facts, with the conditions for every level of every criterion defined in advance. Every vendor named in a review receives a right of reply before publication. The method is public, versioned and reusable by anyone under CC BY-SA 4.0.

Version 1.1 changes. Three fixes from the first review's feedback. Annex B, criterion 4, level 2 gains a third branch, material enforcement with full remediation on record, so a remediated sanction has a correct level instead of falling between them. Criteria 1 and 5 each state their division of labour, so the same fact cannot be argued as counted twice. The scope section states how vendor commitments that reach beyond the EU and EEA are handled. Annex A gains a Status field. No weights, tiers or other level conditions change.

Version 1.2 changes. One extension, prompted by the first European review. The top levels of criteria 1 and 5 and of the money-flow proxy were anchored to the EU and EEA alone, which structurally misread vendors in adequacy jurisdictions: Switzerland is EFTA, not EEA, and carries GDPR-standard protection with no extraterritorial compulsion regime. Where a top level turns on jurisdiction, a jurisdiction holding a valid EU adequacy decision under Article 45 of the GDPR as at the capture date now reads as the EU and EEA for that condition. The extension is named, dated and narrow; adequacy decisions lapse and are reviewed, and the review records the decision in force at capture. No weights, tiers or other level conditions change, and no published review's score moves: both vendors reviewed to date are US-parented, and the clause touches nothing in their scores.

Why a public method

Technology review is a crowded trade, and almost all of it scores the same things: features, price, performance. Those things matter, and this scorecard does not pretend otherwise. Questions that matter most to a European buyer are scored nowhere, because the obvious bodies to score them are paid by the vendors they would score.

Where the data sleeps. Which courts can compel it. What the terms of service permit when nobody is watching. Whether the continent can keep, maintain, defend and support the relationship for ten years without hollowing out its own capability. A buyer weighing these questions today finds opinion dressed as review, or silence.

An independent practice can fill the gap, but only under one condition: the method must be public. A score without a published method is an opinion with a logo. A published method turns every disagreement into an argument about evidence, and an honest reviewer is willing to have that argument. The vendor may dispute a weighting. It cannot dispute its own terms of service, its own enforcement record or its own emissions filing. That asymmetry is the entire design.

A public method compounds. Procurement teams can apply the same eight tests without waiting for a review. Journalists can audit a score in an afternoon. Other reviewers can adopt the framework and argue with it in the open. The scorecard stops being content and becomes infrastructure.

What the scorecard is, and what it is not

The scorecard reviews products and services on which European organisations depend: cloud platforms, productivity software, communications tools, AI services and security tooling. It scores the relationship between the product and European law, European values and European capability.

The scorecard's anchors are EU and EEA instruments, and its scope is the EU and EEA. UK organisations should read criteria 1, 4 and 5 against UK GDPR and ICO enforcement, which diverge from the EU position in ways that change scores materially. A UK annex follows as a method update. Where a vendor's own commitments extend beyond the EU and EEA, for example to EFTA-only Switzerland, the review scores the EU and EEA position and notes the divergence. Where a criterion's top level turns on jurisdiction, a jurisdiction holding a valid EU adequacy decision under Article 45 of the GDPR as at the capture date reads as the EU and EEA for that condition.

Three limits apply, stated here so they are never implied elsewhere.

The scorecard is no security certification and does not test for vulnerabilities. A high score is no guarantee of safety.

The scorecard is no legal advice. It reports what the law requires and what a vendor does, and it does not tell any organisation what to buy.

Scores are dated statements of evidence at a point in time, not predictions. When the evidence changes, the score changes, and the change is logged.

The eight criteria

Eight criteria sit in three blocks. Block A scores the technical relationship. Block B scores what the vendor owes the citizen. Block C scores the decade, because every purchase is also a long-term commitment by the continent.

BLOCK A · THE TECHNICAL RELATIONSHIP

Criterion 01

Residency and jurisdiction

Where data is stored and processed, which legal regimes reach it at rest and in transit, and what that exposure means in practice. A server in Frankfurt that answers to a foreign court is a different proposition from a server in Frankfurt, full stop. The criterion scores where the data sits and what reaches it; who owns the vendor is scored by criterion 5.

Criterion 02

Encryption and key control

Who holds the keys, who can rotate them, and whether the vendor can read customer data without the customer. End-to-end claims are scored against what the architecture and the terms permit. The criterion scores data access governance, not cryptographic strength, and it is no assurance scheme such as SOC 2 or Cyber Essentials.

Criterion 03

Exit terms

What leaving costs, in money, in time and in data. Whether the customer can take a complete copy in open formats, how long the vendor may retain what was surrendered, and whether the contract punishes departure.

BLOCK B · WHAT THE VENDOR OWES THE CITIZEN

Criterion 04

Transparency record

GDPR enforcement and fines, published transparency reports, breach notification history. The criterion scores the record, not the brand reputation that floats above it.

Criterion 05

Ownership and ultimate control

Who can compel access to the data today: the parent company, its government, its courts. The criterion scores the chain of command above the brand on the tin, because that chain, not the brand, decides what happens under pressure. The criterion scores who owns and controls the vendor; where the data sits is scored by criterion 1.

Criterion 06

Citizen protection in the contract

What the terms permit: training AI models on customer data, selling metadata, retention after exit, unilateral change. The test is what the contract allows, not what the marketing promises, because the contract is what the customer signs.

BLOCK C · THE DECADE

Criterion 07

Environmental impact

Power usage effectiveness, renewable share, emissions disclosure under the CSRD where it applies, water use and data centre location. Infrastructure Europe cannot power is infrastructure Europe cannot keep.

Criterion 08

European sustainability

Whether the region, the country and the continent can afford to keep, maintain, defend and support the relationship. The criterion runs through three named proxies, each scored 0 to 4: money flow (where the revenue goes), capability (whether European engineers are employed and trained, and whether the product can be supported without the vendor's home government) and continuity (the vendor's viability on published financials and audited accounts as at the capture date, and its history of withdrawn services). The criterion score is the lowest of the three proxies, reported beside them. Sustainability is a chain, and the weak link governs. The criterion scores what is known as at the capture date. It never scores a projection.

The first three criteria ask what the product does. The next three ask what the vendor owes the citizen. The last two ask what the continent can live with. Together they form the question this scorecard exists to answer. A review never asks whether the technology is good. It asks whether the relationship is one Europe can afford.

How scoring works

Every criterion scores 0 to 4, with fixed definitions. Half marks are not used. A claim without evidence never scores above 1, no matter who makes it. The conditions for every level of every criterion are fixed in advance in Annex B. Two reviewers applying the same evidence to the same rubric should reach the same score, and where they do not, the disagreement has a named address.

SCORE	DEFINITION
4 · Proven	The conditions for level 4 are met, the evidence is public, and independent corroboration exists.
3 · Demonstrated	The conditions for level 3 are met and the evidence is public, without independent corroboration.
2 · Partial	The conditions for level 2 are met in part, or the evidence is incomplete or dated.
1 · Weak	The conditions are largely unmet, or the vendor offers claims without evidence.
0 · No evidence	Nothing checkable supports the standard, or the evidence contradicts it.

The overall score is the simple mean of the eight criteria, reported to one decimal place, with the three block scores published beside it. All eight criteria carry equal weight. If every criterion is a headline, none is. A reader who disagrees with the weights can say so in public, because the weights are published here. That argument is a feature of the method, not a failure of it.

A tier is always displayed with its capture date: Sovereign as at a stated date, never Sovereign alone. A dated score wearing an undated verdict would break the snapshot principle, so the date travels with the label everywhere the tier appears.

Every score is re-examined at least every 12 months, and criteria 1, 4 and 5 are re-checked at each quarterly roll-up, because jurisdiction, enforcement and ownership change fastest. Evidence older than 12 months is flagged stale, and a stale score drops out of the Sovereignty Index until refreshed. No score sits unaltered in the index on aged evidence.

Overall scores map to one of three tiers:

SCORE	TIER	MEANING
3.0 to 4.0	Sovereign	The relationship can be defended to a board, a regulator and a citizen, with the evidence on the table.
2.0 to 2.9	Conditional	The relationship is workable under stated conditions, and the conditions belong in the procurement file.
0.0 to 1.9	Dependent	The relationship depends on the vendor's goodwill, and goodwill is not a control.

Evidence rules

Six rules govern every score.

01 **Facts, not impressions.** A score rests on checkable evidence: terms of service, privacy notices, transparency reports, CSRD filings, enforcement records published by national data protection authorities, the EU Code of Conduct on Data Centre Energy Efficiency, and the product itself where hands-on testing is possible. Every item of evidence carries its capture date.

02 **Anchors in law.** The values block is anchored in the EU Charter of Fundamental Rights, the GDPR, the judgment of the Court of Justice in Schrems II (Case C-311/18), the AI Act and the CSRD. The criterion measures distance from those anchors, not from any reviewer's mood.

03 **Attribution.** Vendor claims are attributed to the vendor. Marketing is evidence of what the vendor wants believed, nothing more.

04 **Right of reply.** Every vendor named in a review receives the draft scores, the evidence behind them and the rubric applied, at least 14 days before publication, with a named contact. Responses are published in full alongside the review, and scores are revised where the evidence requires it. Replies are reproduced verbatim, by licence of their sender. If the vendor does not respond, publication proceeds and the non-response is noted beside the review. If the vendor disputes a score with counter-evidence the reviewer rejects, the counter-evidence is published with the reviewer's reasons, and the reader weighs both.

05 **No test, no claim.** Nothing is scored from a briefing, a demo or a conference call. What has not been evidenced is scored as no evidence.

06 **Corrections.** Errors are corrected openly. The change, its date and its reason are logged at the foot of the review, and a corrected score keeps its history.

Independence

AltLibre takes a small amount of advisory work alongside the research. No vendor commissions, and clients do not influence what is published. That rule extends here: no vendor may pay for a review, a score, a placement or a mention, before or after publication. Any financial relationship with a reviewed vendor is disclosed inside the review itself. Advisory clients are not reviewed while the engagement lasts, and any past engagement is disclosed.

The procedure

Every review carries the same fixed head: the product and version; the category; the dates of evidence capture; the method version used; and the reviewer. Beneath it sits the evidence ledger, one row per criterion, each row naming the evidence and its date. Scores follow, with the rubric conditions quoted beside any score a reader might dispute. Then the vendor's reply in full, then the corrections log.

Quarterly, the published scores roll up into the Sovereignty Index, AltLibre's running picture of where Europe stands. The index carries the roll-up, the capture dates and a short commentary per block for the quarter. Weights stay equal and fixed between method versions. The index is the fixture. The reviews are its evidence.

Reuse

The scorecard is published under CC BY-SA 4.0. Quote it, apply it, translate it. The rubric conditions in Annex B are written to be lifted directly into tender acceptance criteria, and a procurement team can run the eight tests against a shortlisted product in an afternoon. If the method is wrong, the correct response is a better method in public, and AltLibre will publish the argument.

Limits

Honesty about limits is part of the method. AltLibre is a one-person research practice. Testing is desk-bound where hands-on testing is not possible, and a score can inherit the blind spot of its evidence. Scores are snapshots, not verdicts. The method is versioned, and every review states the version it used, so a reader can always tell which rules produced which number.

The first review

Every organisation in Europe already lives inside one of these relationships, usually without knowing which tier it occupies. The first review follows this method within weeks.

The scorecard will be wrong somewhere, and the method is built so the reader can find out where: the evidence ledger, the dates, the weights, all published. A vendor that disputes its score is invited to dispute it with evidence, in public, where the reply will run in full.

> The scorecard does not need to be unarguable. It needs to be arguable, and it is.

References

1. Charter of Fundamental Rights of the European Union, OJ 2012/C 326/02.
2. Regulation (EU) 2016/679 (General Data Protection Regulation).
3. Court of Justice of the EU, Case C-311/18 (Schrems II), judgment of 16 July 2020.
4. European Data Protection Board, Recommendations 01/2020 on measures for transfer impact assessments.
5. Regulation (EU) 2024/1689 (Artificial Intelligence Act).
6. Regulation (EU) 2023/2854 (Data Act).
7. Directive (EU) 2022/2464 (Corporate Sustainability Reporting Directive).
8. Directive (EU) 2022/2555 (NIS2).
9. EU Code of Conduct on Data Centre Energy Efficiency, European Commission Joint Research Centre.
10. Power usage effectiveness, as defined by The Green Grid.
11. Creative Commons Attribution-ShareAlike 4.0 International licence (CC BY-SA 4.0).

Annex A · The review template

Every review uses this fixed structure. Fields in square brackets are completed per review; the bracketed fields are the template, not gaps in it.

Review head

FIELD	ENTRY
Product and version	[name, version]

Vendor and parent	[trading entity, ultimate owner]
Category	[cloud / productivity / communications / AI / security]
Evidence captured	[date range]
Method version	[x.y]
Reviewer	[name]
Status	[draft / at right of reply / published]

Evidence ledger, one row per criterion

CRITERION	EVIDENCE	SOURCE	CAPTURE DATE
1 Residency and jurisdiction	[evidence]	[source]	[date]
2 Encryption and key control	[evidence]	[source]	[date]
3 Exit terms	[evidence]	[source]	[date]
4 Transparency record	[evidence]	[source]	[date]
5 Ownership and ultimate control	[evidence]	[source]	[date]
6 Citizen protection in the contract	[evidence]	[source]	[date]
7 Environmental impact	[evidence]	[source]	[date]
8 European sustainability	[proxy scores: A money flow, B capability, C continuity]	[source]	[date]

Scores, tier always shown with its capture date

CRITERION	SCORE (0 TO 4)

1 Residency and jurisdiction	[n]
2 Encryption and key control	[n]
3 Exit terms	[n]
4 Transparency record	[n]
5 Ownership and ultimate control	[n]
6 Citizen protection in the contract	[n]
7 Environmental impact	[n]
8 European sustainability	[lowest of A, B, C: n]
Block A mean	[n.n]
Block B mean	[n.n]
Block C mean	[n.n]
Overall, mean of eight	[n.n]
Tier, as at [capture date]	Sovereign / Conditional / Dependent

Then, in fixed order: the evidence discussion, criterion by criterion, with rubric conditions quoted beside any score a reader might dispute; the vendor's reply, in full; the corrections log; the disclosures; and the licence note.

License note for each review: "This review is published under CC BY-SA 4.0 by AltLibre. Method: the Sovereignty Scorecard, version [x.y]."

Annex B · The criterion rubric

How to read the rubric. Each criterion carries five levels. A criterion scores the highest level whose conditions all hold; if any condition at a level fails, the score falls to the highest level below whose conditions do hold. Conditions are checkable as at the capture date. They are written so a procurement team can lift them into tender acceptance criteria without the rest of the paper, and so two reviewers holding the same evidence can be forced to the same number.

1 · Residency and jurisdiction

LEVEL	CONDITIONS
4	Data is stored and processed in the EU or EEA, or in a jurisdiction holding a valid EU adequacy decision as at the capture date, and no third-country legal regime reaches it in practice.
3	Data is in the EU or EEA; a third-country parent exists, but the terms document a barrier, such as customer-held keys or an EU-controlled entity, that keeps content unreachable without EU legal process.
2	Data is in the EU or EEA, but a third-country regime reaches it through the parent with no technical or structural barrier; or residency is committed for primary regions only and processing locations are unclear.
1	Residency is claimed but not committed in the contract, or processing relies on third countries with an adequacy decision as the only safeguard.
0	Data is stored or processed in third countries without safeguards, or the vendor will not state locations.

2 · Encryption and key control

LEVEL	CONDITIONS
4	The customer holds the keys and the vendor has no technical route to content, as documented and enforced.
3	The vendor holds keys with documented barriers and customer-controlled rotation; any vendor access follows a documented, customer-visible process.
2	Encryption at rest and in transit with vendor-held keys; no customer rotation and no access transparency.
1	Encryption is claimed, the scope is unclear and keys are vendor-held with no customer control.
0	No encryption, or the vendor cannot state the model.

3 · Exit terms

LEVEL	CONDITIONS
4	Complete export in open formats at no charge; termination on 30 days' notice or less; no retention beyond a short stated deletion period; no exit fee.
3	Full export available with a bounded fee or retention period that does not materially burden departure.
2	Export available but partial, in proprietary or lossy formats, or with fees and retention that impose real cost.
1	Export impractical, or the contract is silent on data return.
0	The contract bars export or grants the vendor continuing rights over surrendered data.

4 · Transparency record

LEVEL	CONDITIONS
4	No material enforcement; transparency reports published on a stated cycle; breach notification history prompt and complete.
3	No material enforcement; limited published transparency.
2	Enforcement or audit findings with partial remediation, material enforcement with full remediation on record, or gaps in the breach notification record.
1	Material enforcement without remediation, or repeated findings across periods.
0	Serious enforcement, systematic failure to notify, or no record obtainable.

5 · Ownership and ultimate control

LEVEL	CONDITIONS
4	EU or EEA ownership, or ownership in a jurisdiction holding a valid EU adequacy decision as at the capture date, with no third-country parent and no third-country compulsion pathway.
3	EU or EEA headquarters with third-country investment below control and documented governance barriers.
2	Third-country parent with structural ring-fencing of EU data, such as EU-controlled keys or entity, or ownership disclosure that is incomplete.
1	Third-country parent with a clear compulsion pathway and no ring-fencing.
0	Ownership undisclosed, or control sits in a jurisdiction with broad extraterritorial powers and no mitigations.

6 · Citizen protection in the contract

LEVEL	CONDITIONS
4	The contract bars AI training on customer data and its sale, retains nothing after exit beyond a short stated period, and requires notice and customer acceptance for material change.
3	AI training and sale barred by default, with narrow opt-in exceptions bounded in the contract.
2	The contract is silent on AI training, permits secondary use under a broad improvement clause, or retains data beyond exit for an unbounded purpose.
1	AI training or secondary use proceeds without a meaningful opt-out, or material change is unilateral.
0	The contract grants the vendor broad rights over customer data, including after exit.

7 · Environmental impact

LEVEL	CONDITIONS
4	Power usage effectiveness at or below 1.2, with a majority renewable supply under documented procurement and public disclosure in place.
3	Power usage effectiveness at or below 1.4, or disclosure that is public but partial on energy sources.
2	Partial disclosure, or power usage effectiveness above 1.4 without a stated improvement plan.
1	No public disclosure; participation in the EU Code of Conduct on Data Centre Energy Efficiency or an equivalent commitment only.
0	No disclosure and no efficiency commitments.

8 · European sustainability

Criterion 8 runs through three proxies, each scored on the same 0 to 4 scale. The criterion score is the lowest of the three, reported beside them.

LEVEL	PROXY A · MONEY FLOW
4	Revenue is retained and taxed in the EU or EEA, or in a jurisdiction holding a valid EU adequacy decision as at the capture date.
3	An EU entity retains the majority of revenue, with limited repatriation.
2	An EU entity exists but the majority of revenue leaves the EU and EEA.
1	A nominal EU presence exists for sales only.
0	No EU or EEA entity.

LEVEL	PROXY B · CAPABILITY
4	EU or EEA engineers operate and maintain the product, with documented skills and training investment.
3	An EU engineering presence exists, but core updates depend on the parent organisation.
2	EU staff provide support only, with all engineering elsewhere.
1	No EU or EEA technical employment is checkable.
0	No EU or EEA presence of any kind.

LEVEL	PROXY C · CONTINUITY
4	The latest published or audited accounts show a viable business, and the vendor has withdrawn no service in the category without supported migration.
3	Published accounts show a viable business; no withdrawals, or audited accounts unavailable.
2	The business is viable on published figures but unproven at scale, or has one withdrawal in the category with supported migration.
1	Published figures show financial stress, or repeated withdrawals with unsupported migration.
0	Insolvency proceedings, abrupt withdrawal, or no accounts obtainable.

All three proxies read the evidence as at the capture date. Published accounts are facts; projections are not scored.

ADRIAN HOLLISTER · ALTLIBRE

The Sovereignty Scorecard, the method · Version 1.2, 1 October 2026.
Published under CC BY-SA 4.0. Free to quote, republish and translate.
altlibre.com