BHI
Black Hole Index
STRUCTURAL LOCK-IN MEASUREMENT · B = CAPTURE / ESCAPE

blackholeindex.com
ivan.savich@blackholeindex.com
ORCID 0009-0003-9542-7022 · SSRN 6555158

Information provided ahead of the proposed decision: a structural profile of Microsoft's integrated business-software stack

A frame-first measurement — boundary frozen on 22 August 2026 before any evidence was gathered, the S2 package whose manifest fixes its 1.0.8 bytes publicly hash-fixed and independently timestamped on 29 August 2026, its successor S4 (boundary 1.0.9, profile 1.0.7) on 31 August 2026 (see §7.2), amendments disclosed — of whether integration across the products creates structural dependence greater than the sum of the products, read from Microsoft's own documentation, and the named substitutes' where substitutability and the network standards test require it, for one stipulated UK enterprise, using the Black Hole Index.

Submitted by: Ivan Savich, Creator of the Black Hole Index (independent researcher)

Case: SMS investigation into Microsoft's business software ecosystem (Investigation Notice, 14 May 2026)

Channel: the case-team mailbox named on the case page (business.software.sms@cma.gov.uk); unsolicited, ahead of the proposed decision indicated for October 2026

Declaration: I have no commercial relationship with, and receive no funding from, Microsoft or any party with an interest in this investigation; no party with an interest in it has funded, commissioned or reviewed the record. All source documents relied on as evidence are public; every quotation in the scored rows is byte-verified against a retained capture, whose per-file digests are in the S2 and S4 manifests (§1.3, §7.2), with the exceptions in §1.3 and the non-record facts declared in §7.2. This document is non-confidential and may be published in full.

Record: measurement boundary v1.0.8 (revised 1.0.9 on 30 August 2026, wording only), rules delta v1.7, evidence journal v1.0.2, structural profile v1.0 (rev. 1.0.1–1.0.7, provenance and limitation clauses), financial-exit register v1.3 — these versions fixed 29 August 2026 (boundary v1.0 frozen 22 August 2026), public snapshots S2 and S4 (S2 fixes the 1.0.8 and profile v1.0 bytes; S4, frozen 31 August 2026, fixes the 1.0.9 and profile v1.0.7 bytes, the 30 August revisions dated, archived and hashed inside it)

1. Introduction and standing

1.1. I provide this information to the case team for the CMA's SMS investigation into Microsoft's business software ecosystem (Investigation Notice of 14 May 2026), ahead of the proposed decision the indicative timetable places in October 2026. It is unsolicited: the invitation to comment closed on 4 June 2026, and I make no request beyond that the record be available to the case team and, if useful, cited. I am an independent researcher and the Creator of the Black Hole Index (BHI), an open measurement instrument for structural lock-in on digital platforms, published at blackholeindex.com — its methodology, rubric and quarterly scores open under CC BY 4.0; the CMA published my response to its Apple steering consultation among the named responses on 14 August 2026 (publication, not endorsement).

1.2. Nothing in this document is a view on whether Microsoft has, or should be designated as having, strategic market status; that determination is the CMA's. What follows is a measurement record: a structural profile of Microsoft's integrated business-software stack as consumed by one stipulated UK enterprise, read from first-party documentation — Microsoft's, and the named substitutes' where substitutability and the network standards test require it; six of eight values read from documentation, two stipulated — under a frame frozen before any evidence was gathered (22 August 2026) and hash-fixed

publicly on 29 August 2026 in snapshot S2 (the public digest of that fixation is that of the S2 package, whose manifest fixes the boundary's 1.0.8 bytes; see §3.6 and §7.2), with the post-hoc amendments and post-scoring reading rules disclosed in §3.6, §4.3 and §6. It asks one question the Notice's ecosystem framing raises directly: whether integration across the products creates structural dependence greater than the sum of the products taken separately.

1.3. I have no commercial relationship with, and receive no funding from, Microsoft or any party with an interest in this investigation; no party with an interest in it has funded, commissioned or reviewed the record. Every source document relied on as evidence is public. Every quotation in the scored rows is byte-verified against a retained capture, except four confirmed verbatim by a browser-class fetch where the page returned empty or as a script shell — a recorded channel difference. Every frozen version is hash-pinned with its superseded bytes archived, with the declared exceptions recorded in the profile's limitations, provenance and revision sections and in the boundary's revision history, together with one further exception visible only in the comparability register's own revision history (its superseded 0.1/0.2/0.3 bytes, not archived; the comparability register, the structural profile and the standalone-anchor matrix carry no SHA pin in the enforcement script, their S2- and S4-state bytes being fixed by those packages' manifests) (boundary bytes before 1.0.3, declared in boundary revision 1.0.4; in-place corrections within a revision on 26 and 29 August; several register builds during the 29 August audit passes that are not archived but are hash-regenerable; one in-place edit of the profile's provenance line on 30 August, recorded the same day as profile revision 1.0.1 with the snapshot bytes archived; the checker's bytes before boundary 1.0.7 and its intermediate states within 1.0.8, re-archived at each pin rewrite; and the quote verifier's bytes before its 29 August normalisation change); the frozen packages were fixed publicly and independently timestamped (RFC 3161) on 29 August and 31 August 2026 (§3.6).

2. Summary of the record

2.1. Eight of the instrument's eleven parameters carry a value for the stipulated customer. Six are read from first-party documentation: Microsoft's, and the named substitutes' where substitutability and the network standards test require it; three of the six (data, memory and action) are tagged part frame-read in the record's provenance column and say so; the network row's frame-read legs do not carry its step. Two are stipulated by the frame (§2.3). Three — closeness, human fallback and momentum — closed as Not Scorable, and no composite score is computed (§7); two declared weaknesses the reader should carry — the network reading's non-record fact and substitutability's unanswered "at integration level" qualifier — are stated at §7.2. The record therefore prints parameter values and their decomposition, not a single number.

2.2. On the capture side the integration thesis holds where it is decidable — at full perimeter; at investigation scope on data, memory and network, the action increment falling to zero there; at activity scope only the data (+2) and network (+1) increments survive (§4.2): the documented cross-product mechanisms add to the readings beyond what any component supplies alone — the data graph (item and people insights across documents, SharePoint and Teams), cross-plane correlation in the security estate and the collaboration mesh to data; the graph, the security signal and the collaboration mesh to memory (the AI-grounding leg part frame-read); one distribution unit and one Graph API endpoint across hosts to network — and conditionally to the action reading. On the escape side it fails on the reading of record: portability and substitutability carry no integration increment — the documented conditions attach surfaces, roles, sequencing and federation to a component exit without removing a substitute or closing an export path. A reading registered after scoring and since confirmed (a −1 on each escape parameter

where a documented condition attaches to the exit) is recorded beside the record and co-equal in evidential weight, printed in the same cell for space; on it the thesis would hold on the escape side too.

2.3. Two of the eight values are stipulations of the frame, not measurements — process (10, on the sole-production-host clause of the reference customer, Not Scorable under 1.0.3, the boundary as it stood when scoring began) and organisational depth (7, on its tenure, deployment and the existence — not the terms — of a single enterprise agreement, with one clause read as near-certain rather than deductive, a declared deviation) — and are printed as such. Every capture basis except network's is read on the cloud platform component — for memory, action and process it is set by that component alone — which lies outside the Notice's five digital activities; for that reason every scored row prints two further columns, at activity scope (the Notice's activities) and at investigation scope (activities plus the AI assistant layer), and the divergence arises on the three capture rows the cloud platform sets alone, so a reader confined to the Notice can take the activity-scope column.

3. Method: the frame, fixed before scoring, with its amendments disclosed

3.1. **Object and membership test.** Microsoft's integrated business-software stack, defined not as a product list but by a membership test: a component is inside the perimeter only if (a) the stipulated customer consumes it and (b) its own first-party documentation shows tenancy in the stack's identity plane — the component authenticates its users, devices or workloads against the organisation's Entra tenant, as the enterprise posture the vendor's current deployment documentation prescribes for an organisation of the Reference Customer's shape. Ten components pass (§5, table): productivity and collaboration suite; real-time communication; workforce identity and access; endpoint management; security estate; client operating system; server operating system; relational database engine; cloud platform; AI assistant layer. GitHub, Dynamics 365, Fabric and Power Platform are excluded with grounds stated, as are consumer services, hardware, gaming, advertising and LinkedIn; the declared consequence is that a real UK enterprise that does consume Dynamics 365 or Power Platform is a more bound customer than the one measured, so this perimeter is a conservative understatement relative to a customer consuming the wider estate, and no result from it may be quoted as describing that wider customer.

3.2. **Reference Customer.** One stipulated UK private enterprise: 1,200 staff (about 150 engineering and platform, about 40 IT operations and end-user computing), a digital product delivered as a service, six years on the cloud-integrated form of the stack, hybrid identity (on-premises Active Directory synchronised to an Entra tenant that is the cloud identity root), about 1,300 managed devices on the vendor client OS, the AI assistant deployed organisation-wide, about 120 TB of data across the stack, one enterprise agreement. One organisation, not a range, because a range lets the object be scored against its most flattering point; every anchor is applied to this organisation; three stipulations — UK geography, the productivity suite consumed, and the proprietary relational engine — deliberately make it harder to exit than the reference customer of the earlier BHI Cloud Core determination (four cloud platforms, snapshot S1), and the record says so.

3.3. **Integration-links register.** "Integration" is given measurable form as eight documented links (IL1–IL8), in the boundary's own names: the shared identity root; suite licensing and administration; the data graph; AI grounding (on organisational content held in the suite, communication and identity components, as narrowed in boundary 1.0.3); the cross-product security signal; device–identity binding; directory–cloud coupling; and the collaboration mesh. Each scored parameter's reading is decomposed into a standalone basis and an integration increment, and each increment is attributed to named links with their measurement provenance: provider-documented (P), read from the frame (R) or both (R+P; described in §4 as "part frame-read"). The rows are not independent: the same documented mechanisms

feed several parameters, and the shared identity root — which is also this perimeter's membership test — is cited in seven of the eight scored rows; the record names every correlated input.

3.4. **Exit criterion.** A completed stack exit is defined once, in advance: within twelve months, with the customer's existing engineering, platform and IT-operations staff plus contracted assistance not exceeding 25% of that combined headcount, the identity root is reconstructed off the vendor's plane, the device fleet re-enrolled into a non-vendor management plane and, to the extent the client OS is replaced, re-imaged, the collaboration estate migrated with business continuity, the server and relational estate on non-vendor engines or the customer's own infrastructure, the cloud workloads satisfying the Cloud Core exit definition (as defined in that determination, snapshot S1), a functioning security estate in place, and no business function withdrawn. An export capability is not an exit.

3.5. **Evidence and price rules.** Microsoft's own documentation is the evidence on the capture side; substitutability is searched symmetrically — the same question put to each component and to its named best-documented substitute before any answer is used, so that "no documented equivalent" claims are tested on both sides, with the three components whose substitute sides were not recoverable declared at §7.2; the network row's standards test reads the substitutes' own documentation for that test only; portability is searched symmetrically across the components, on the vendor's documentation. The evidence window is 23 August to 19 September 2026 (opened two days early by a dated clarification, boundary 1.0.1, before any evidence was gathered); the rows were read on 26 and 29 August 2026, every retrieval date is per-row in the record, and the determination closed on 29 August 2026 and is frozen — the window's later cut-off bounds source publication dates, not this record. Vendor competitive claims and marketing comparisons are never sufficient on their own; price and licensing terms enter no structural parameter and are recorded separately (§6). Expected determinability was declared before scoring: closeness and human fallback were expected to be Not Scorable from provider documentation and momentum for want of a common unit of capability change — and each was. The scoring order and the single evaluator's anchoring exposure were declared in the boundary before any row was scored.

3.6. **Fixation, and the two substantive amendments after scoring began (with every other post-scoring revision listed).** Boundary 1.0.4 (26 August 2026) added the sole-production-host clause of the reference customer after data, memory and action had been scored; a first writing of process at 10 had been withdrawn for want of it and was re-entered under 1.0.4 (direction: harder to exit), and the record prints process forked for that reason (§4). The protocol treats a boundary change as a new measurement; this record treats 1.0.4 as one measurement because the rows scored before it were re-read under it without movement — data and action on 26 August, memory on 29 August during the audit — and process, the only value existing solely under 1.0.4, is printed forked. Boundary 1.0.6 fixed the countable set of the financial-exit register (§6) after the searches had run — a post-hoc definition, disclosed as such. The remaining post-scoring revisions (1.0.5, 1.0.7, 1.0.8) changed no stipulation, and the pre-scoring clarifications 1.0.1 and 1.0.3 are noted at §3.5 and §3.3 (1.0.2 corrected the date attribution of the 1.0.1 row); revision 1.0.9 of 30 August 2026 names the document in the anchoring-exposure declaration by its title (boundary §11; quoted at §7.1 below), dated, archived and hashed outside snapshot S2, which fixes the 1.0.8 bytes, and inside snapshot S4. Measurement boundary v1.0.8, rules delta v1.7, evidence journal v1.0.2 (frozen v1.0 29 August 2026; revised 1.0.1/1.0.2 the same day, status and declarations only — the 1.0.2 bytes are the ones in S2), the structural profile v1.0 (revised 1.0.1–1.0.7 on 30 August 2026, provenance and limitation clauses — those seven revisions dated and archived outside S2, each superseded state hashed, S2 fixing the profile's v1.0 bytes and S4 the 1.0.7 bytes) and the financial-exit register v1.3 are inside the package whose SHA-256 is fixed in the public snapshot record S2 (blackholeindex.com/research-snapshots.txt) and RFC 3161-timestamped by two independent authorities on 29 August 2026; the record's state as despatched — boundary 1.0.9, profile 1.0.7, every retained su-

perseded version archived, with the declared exceptions at §1.3 and §7.2 — is fixed the same way as snapshot S4, timestamped by the same two authorities on 31 August 2026; the packages themselves are private pending quality assurance — not publicly published, but available to the case team on request (§8) — and the per-document digests are in their manifests. Two reading rules registered after scoring (§4.3) are disclosed as such; no score, clause coding or §9 count changed after the determination's freeze of 29 August 2026 (§3.5; "the freeze" below refers to it); what did change after it is which registered reading and which of two computed counts is the record, and the boundary's own definition of the countable set (1.0.6, adopted after the freeze) (§4.3, §6, §6.3).

4. Results: the structural profile

4.1. Parameters are scored on the published 0–10 anchors; "basis + increment" separates what the component supplies alone from what the documented integration adds. Activity scope reads only the components inside the Notice's five digital activities; investigation scope adds the AI assistant layer. Both scope columns are derived from the same rows by striking the cloud-platform legs — and, at activity scope, the AI-assistant-layer legs — one by one, not re-scored (organisational depth is stipulated at organisation level and is unchanged by either striking); for the data row the record states that the anchor distance was not re-assessed once the cloud platform is struck. Not Scorable means the parameter cannot be read for this object — from provider documentation (closeness, human fallback) or for want of a common unit (momentum) — not that it is zero. The capture and escape values are not combined; no ratio computed from this table is a BHI result. Substitutability is not a cohort-relative score: it is a statement about the named substitute set as documented in the evidence window (§3.5) and may not be quoted as if it came from a scored cohort. The named substitutes, per component: suite — Google Workspace, LibreOffice; real-time communication — Slack, Zoom (Cisco Webex read for the network standards test); identity — Okta; endpoint management — Omnissa Workspace ONE UEM, Jamf Pro; security estate — CrowdStrike Falcon; client OS — macOS, Ubuntu; server OS — Samba AD DC, Red Hat Identity Management (the component is excluded from the basis); database engine — PostgreSQL, and Babelfish for Aurora PostgreSQL; cloud platform — AWS, Google Cloud; AI assistant — Gemini for Google Workspace, ChatGPT Enterprise.

PARAMETER	READING OF RECORD (FULL TEN-COMPONENT PERIMETER)	BASIS + INTEGRATION INCREMENT	ACTIVITY SCOPE	INVESTIGATION SCOPE
Data capture (d)	9 (8 or 10 on the declared reader alternatives)	7 + 2 (provider-documented: the data graph — item and people insights, the cross-product security signal, the collaboration mesh; part frame-read: the identity root, AI grounding, device–identity binding, directory–cloud coupling)	9 (+2)	9 (+2)

PARAMETER	READING OF RECORD (FULL TEN-COMPONENT PERIMETER)	BASIS + INTEGRATION INCREMENT	ACTIVITY SCOPE	INVESTIGATION SCOPE
Memory capture (m)	9 (8–9 on the confirmed modal reading)	8 + 1 (provider-documented: the data graph, the cross-product security signal, the collaboration mesh; part frame-read: AI grounding, the identity root, device–identity binding, directory–cloud coupling)	7–8 (+0)	8–9 (+1)
Action capture (a)	9 conditional; 8 otherwise	8 + 1, conditional on two unstipulated configurations (security automation level; password-hash synchronisation); every leg part frame-read (the identity root, the cross-product security signal, device–identity binding, directory–cloud coupling)	(+0) — 7 cond. / 5–7	(+0) — 7 cond. / 5–7
Process capture (p)	10 under boundary 1.0.4; Not Scorable under 1.0.3, the boundary as it stood when scoring began	10 + 0 (ceiling; a deduction from the 1.0.4 sole-production-host clause, carrying no information about the vendor)	Not Scorable	Not Scorable
Network capture (n)	6 (8 on the standards-authorship reading; 9 on the gated-integration reading)	5 + 1 (provider-documented: one distribution unit; one Graph API endpoint across hosts — the distribution leg is filed under a link the register names only administratively, recorded as a declared stretch; part frame-read: the identity root, AI grounding)	6 (+1)	6 (+1)
Closeness (c)	Not Scorable	customer-state parameter; not in provider documentation	—	—
Portability (x)	5 (4 on the −1 reading)	5 + 0 (documented conditions attach surfaces, roles and sequencing without closing an export path)	5	5

PARAMETER	READING OF RECORD (FULL TEN-COMPONENT PERIMETER)	BASIS + INTEGRATION INCREMENT	ACTIVITY SCOPE	INVESTIGATION SCOPE
Substitutability (s)	5 (4 on the −1 reading)	5 + 0 (documented conditions at zero; three links withdrawn). Modal reading 5 on five of six voting components (two asymmetric components do not vote; one not adjudicated; the server OS excluded from the basis)	5	5
Human fallback (h)	Not Scorable	customer-state parameter	—	—
Momentum (t)	Not Scorable	no common unit of capability change across release channels	—	—
Organisational depth (o)	7, stipulated	7 + 0 (tenure, deployment and the existence — not the terms — of the single enterprise agreement are the frame's own text; anchor 7's "training" clause read as near-certain rather than deductive — a declared deviation recorded in the evidence journal)	7	7

4.2. **Where the integration test holds and where it fails.** At full perimeter the thesis holds on the provider-documented legs of data (+2), memory (+1) and network (+1), and conditionally on action (+1). At investigation scope data (+2), memory (+1) and network (+1) hold and the action increment is zero. At activity scope the memory and action increments are both zero — memory's because its surviving legs are borne by the cloud platform and the AI assistant layer (the mesh leg in that row being the assistant's meeting output), action's because the leg that carries it is gated by the same unstipulated configuration that sets the basis once the cloud platform is struck — so within the Notice's five activities the thesis holds on data (+2) and network (+1) only. It is undecidable on process, which the frame places at the ceiling. It fails on both scored escape-side rows (portability and substitutability; human fallback is Not Scorable and untested) and on organisational depth, whose zero is set by the frame: every positive increment is capture-side, and every escape-side increment is zero on the reading of record. In the component-exit table behind these rows, nine of ten components show a documented condition attached to another component on exit, and in one — workforce identity and access — the tenant object is retained by the perimeter's own membership criterion and the identity root is frontable, with replacement not evidenced on the record either way. That row is where this record leaves the thesis untested; customer-state evidence, or a symmetric substitute-side search recovered through a readable channel — a first-party statement that a substitute replaces the workforce directory as root — is what would decide it.

4.3. Two reading rules were registered after scoring and are disclosed as such: a modal-versus-extreme reading of each standalone anchor, and a rule that registers, beside each escape parameter, a −1 reading where a documented condition attaches to the exit. Both were adopted as co-equal readings by the record's author on 29 August 2026, after scoring, under my general instruction and without an itemised

decision on each, and they remain open to reversal by me — the record says so; they are printed beside the reading of record, which was not re-scored. The range the journal rows pre-committed to, with the double count declared, is portability 5 to 2 and substitutability 5 to 3 — stated as a declared bound and not adopted; without the double count the range runs from the reading of record to the lock-in-favouring extreme reading: portability 5 to 3, substitutability 5 to 4 — portability's extreme is set by the AI assistant layer alone and the activity-scope extreme is 5; substitutability's is set by the cloud platform, which lies outside the Notice's activity list (§5). The modal reading is adjudicable on three rows only (memory, portability, substitutability); on data, action, process, network and organisational depth fewer than six components carry a standalone anchor, so the modal reading is printed as not adjudicated and the basis rests on the strongest single component — except organisational depth, whose basis is the frame's own stipulation and where no component reading exists. On the capture side the same rules move only memory (9 to 8–9); the data row carries declared reader alternatives at 8 and at 10, the action row a declared downward alternative — a basis of 7 for a reader who confines the rubric to agentic planning — and the network row standing alternatives of 8 and 9.

5. Alignment with the Notice's digital activities

5.1. The perimeter is fixed by the membership test and the stipulated customer, not by the Notice; the alignment is convenient and is not a ground. Read against paragraphs 5, 8 and 9 of the Notice, the mapping is close but not total.

COMPONENT (FUNCTIONAL DEFINITION)	VENDOR MAPPING	NOTICE DIGITAL ACTIVITY
C1 Productivity and collaboration suite	Microsoft 365 (Word/Excel/PowerPoint, Exchange Online, SharePoint, OneDrive)	(a) Productivity Software Suite
C2 Real-time communication	Microsoft Teams	(a) Productivity Software Suite
C3 Workforce identity and access	Microsoft Entra ID	(e) Security Software
C4 Endpoint management	Microsoft Intune	(e) Security Software
C5 Security estate	Microsoft Defender family	(e) Security Software
C6 Client operating system	Windows 11	(b) PC Operating System
C7 Server operating system (incl. Active Directory Domain Services)	Windows Server	(c) Server Operating System; (e) for AD-related services
C8 Relational database engine	Microsoft SQL Server	(d) RDBMS

COMPONENT (FUNCTIONAL DEFINITION)	VENDOR MAPPING	NOTICE DIGITAL ACTIVITY
C9 Cloud platform	Microsoft Azure	Outside the five activities (addressed by the CMA's separate cloud actions)
C10 AI assistant layer	Microsoft 365 Copilot	Copilot features in the apps and/or the Microsoft 365 Copilot App are within the investigation (Notice paragraph 8(a)); whether AI products sit inside the activity definitions is itself under investigation (paragraph 9); included here by this frame's own stipulation

5.2. The column to quote from §4 is the activity-scope column, the reading a reader bounded by the Notice should use (§2.3). The full-perimeter column is the reading of record and the investigation-scope column is the profile's own headline; both are printed beside it for the reasons stated, and neither is the column to quote here. Nothing here estimates the Notice's activities as such.

6. The financial exit quantity, kept separate

6.1. Price and licensing terms move no structural parameter. This is the instrument's design thesis, and it answers a foreseeable objection: the mechanism the CMA's own paper of 31 March 2026 foregrounds is licensing practice. The paper records that the CMA has not identified material progress on the concerns with Microsoft's licensing practices, which it describes as reducing competition in cloud, and states that this was a factor in its decision to prioritise this investigation. A reader will therefore ask why the headline mechanism is outside the index. Price terms can be repriced by a decision; structure has to be rebuilt.

6.2. They are counted separately as the number of clauses in public terms that would bind the stipulated exit for a holder of the instrument each clause belongs to (on the vendor side, one enterprise agreement — including clauses whose enrolment type, hybrid-benefit use and consumption commitment the frame does not fix), in each direction. The vendor side records at least nine binding clauses; on the form test of record (does the sentence, on its own words, grant or restrict?) at least six are coded as raising the cost of the stipulated exit and at least three as lowering it (lower bounds: one terms document returned an access error on every export and archive endpoint tried under the register's pre-fixed source class — downloadable hashed documents; the browser channel the record uses elsewhere (§1.3) is excluded by that specification — UK-specific enterprise-agreement terms are not public, and one term class was found in no reachable vendor document); on the alternative effect test (a right conditioned on a vendor-side purchase, such as active Software Assurance, is read as raising the cost of exit) the vendor side records at least eight raising and at least one lowering. The register records the same quantity for two substitute platforms under the same rule — AWS 8 raising / 1 lowering (effect test 9 / 0), Google Cloud 10 / 0 (10 / 0) — and draws no comparison across sides — the vendor side is a lower bound, and the substitutes' bundling and private-pricing terms were not reachable or not public, and on one substitute no AI add-on dependency of the class was found after a defined search; nor does this document. One clause class on the substitute side was coded against the specification's pre-registered expectation, and the register says so.

6.3. Both tests are printed because three grammatically granting sentences — two on the vendor side, one on a substitute's — code opposite ways under the two tests. That split must be seen before the raising/lowering split is read as an effect measure. The binding yes/no coding was made in the clause search under the pre-fixed rule; the direction coding was added afterwards by the record's author, and

both codings, with their dates, are in the record. Which of the two tests is the test of record was fixed after both counts were known (a disclosed amendment to the register's specification); the form test chosen is the one closer to the pre-fixed coding rule and the vendor-favourable one — though that rule was ambiguous for exactly the three sentences at issue, and the grant-wins tie-break that constitutes the form test was itself made after they were on the page; the register's specification says so. The count is of clauses available in public terms, not of clauses known to be in force for the stipulated customer. The substitutes' raising clauses are of two kinds — conditions on running Microsoft products there which reference the customer's Microsoft agreements, and the substitutes' own commitment constructs — and the post-hoc countable-set reading (§3.6) added more raising clauses on the substitute side than on the vendor side. This quantity is a count of terms, not a structural reading, and is not combined with §4.

7. Methodology status and limitations

7.1. The eight scored rows were produced by a single evaluator — the record's author, who scored every row and recorded every step of the determination — working under my instruction as principal investigator. The per-component evidence searches and the §6 clause searches were run separately from the scoring, under the fixed question sets and budget truncations the journal declares row by row and, for the §6 searches, the term classes fixed in the register's specification before any source was read, with the search records and truncations printed in the register. The record names both roles wherever my instruction bears on a decision, and the three post-scoring choices (the two reading rules taken together, §4.3; the direction test of record and the countable-set reading, §6.3 and §3.6) are recorded as the author's under my general, non-itemised instruction and as reversible by me. Before despatch I have read this document in full and I am responsible for every value and statement in it. The §6 searches' two-stage coding is described at §6.3; captures and byte-verification were run as recorded procedures. The evaluator's anchoring exposure is declared in the boundary and is the strongest yet: the evaluator has read, at first hand and in full, the CMA's paper *Actions on cloud and business software through the UK Digital Markets Competition Regime* of 31 March 2026 (boundary revision 1.0.9 names it so; the 1.0.8 text in snapshot S2 identified it by month and subject rather than by title — the same document); has seen the instrument's quarterly release row for Microsoft (a different object — the conglomerate "Azure + M365 + GitHub" entity — and not an input to this record); and completed the Cloud Core determination containing an Azure column. The exposure was realised: the first writings of six rows each printed the more lock-in-favouring value and were withdrawn or moved away from it on internal adversarial review before the freeze — network lowered from 8 to 6 and organisational depth from 8 to 7, closeness withdrawn to Not Scorable, process withdrawn and re-entered at 10 only after boundary amendment 1.0.4, and the two escape rows raised from 4 to 5. In a second form recorded by the journal, four rows (data, network, substitutability, organisational depth) held their first score across a replaced derivation on the first review pass; three were later moved off that first score (network 8 to 6, organisational depth 8 to 7, substitutability 4 to 5), and data held 9 across every replaced derivation of that row. The instrument's inter-rater reliability is not established and remains its binding constraint; the provenance discipline and the instrument's validation are two different levels and are stated separately.

7.2. Provider documentation cannot carry closeness or human fallback, and no common unit carries momentum; the record prints Not Scorable rather than an estimate, and no composite score is computed below eleven of eleven parameters. Two values are stipulations of the frame (§2.3). Two configurations that condition the action reading are not stipulated and are printed as conditions. Three components' substitute sides were not recoverable by the quotation verifier: endpoint management and the security estate rest on the vendor's own description of the substitute class, and the server operating system is excluded from the basis. Two declared weaknesses the record singles out for a CMA reader (pointed to at

§2.1): network's reading moved from a first writing of 8 to 6 on a non-record fact (the status of ISO/IEC 29500, cited as public knowledge), with Strict and Transitional conformance classes undistinguished — that the default write format is Transitional is itself a non-record fact, stated as such — a declared divergence from the Cloud Core AWS network precedent (§3.2); and substitutability's rubric qualifier "at integration level" is unanswered: no stack-level substitute was searched for or named, so the value is a per-component reading, and the two root components (identity and the server operating system) do not set it. One rubric question is open (whether the closeness rubric applies to an object containing the operating system); per-component search truncations are declared row by row in the journal. For four rows (data, memory, action, process) the fixed question set is the text printed in the row and no separately timestamped question file exists — a declared gap in the pre-registration evidence; the AI-grounding link is a stipulation of the frame and is read as part frame-derived wherever it carries a leg. Momentum's probe covered three of the ten components, a declared truncation; closeness's refusal rests on the top anchor, with the middle anchors not walked; human fallback is bounded above by deduction from the process value — anchor 10 is excluded — and has no point value. The boundary's bytes were first retained at revision 1.0.3 (26 August 2026); the 22 August frozen state is evidenced by the boundary's dated revision history, not by an independently fixed digest, and the public digests are those of the S2 package (manifest: boundary 1.0.8, profile v1.0) and the S4 package (manifest: boundary 1.0.9, profile 1.0.7); the packages are available to the case team on request (§8). One non-binding, unprinted register quotation was re-stitched with an ellipsis, and one further non-binding quotation was brought within the verifier by a whitespace-normalisation change to the verifier itself, both after the freeze; the re-stitched quotation's superseded evidence-file bytes are archived and hashed, and the verifier's pre-change bytes were not retained — a declared exception; from that point every binding and non-binding quotation verifies. The two post-scoring reading rules (§4.3) and the accidental overwrite and hash-verified restoration of the pinned register during audit stand as limitations of the record.

7.3. The measurement reads published documentation, not behaviour: whether the stipulated customer could in practice complete the §3.4 exit within twelve months is a question for customer-state evidence, which the CMA's qualitative customer research is placed to gather and this record is not. A determination under this boundary does not establish market power, dominance, gatekeeper or SMS status, abuse, unfair conduct, infringement, consumer harm or legal liability.

8. Availability

8.1. The frozen boundary, evidence journal, structural profile, financial-exit register and its specification, the standalone-anchor matrix, the comparability register (membership test and link confirmations, 18 of 18 confirmed before any score) and the rules delta, with every quotation, capture hash and reading rule, are available to the case team on request; the package digests are public (S2 and S4), with any post-despatch revision to be dated, archived and hashed outside them (§3.6), and the per-document digests travel with the packages. Should the proposed decision be published for consultation as scheduled, I intend to respond to it on the same record. On despatch of this document the boundary's and the profile's status line "not yet put to it" will be amended by a further dated revision, so that the record and the fact agree. I remain available to answer questions on the method.

Ivan Savich

CREATOR, BLACK HOLE INDEX · INDEPENDENT RESEARCHER

ivan.savich@blackholeindex.com · blackholeindex.com · ORCID 0009-0003-9542-7022

References

Competition and Markets Authority (2026). SMS investigation into Microsoft's business software ecosystem: Investigation Notice, 14 May 2026; case page last updated 23 June 2026, gov.uk/cma-cases/microsofts-business-software-ecosystem

Black Hole Index (2026). Integrated Stack — measurement boundary v1.0.8 (revised 1.0.9 on 30 August 2026, wording only), rules delta v1.7, evidence journal v1.0.2, structural profile v1.0 (rev. 1.0.1–1.0.7, provenance and limitation clauses), financial-exit register v1.3 (these versions fixed 29 August 2026; boundary v1.0 frozen 22 August 2026). Public digests: snapshots S2 and S4 (S2 fixes the 1.0.8 and profile v1.0 bytes; S4, frozen 31 August 2026, fixes the 1.0.9 and profile v1.0.7 bytes), blackholeindex.com/research-snapshots.txt

Black Hole Index (2026). BHI Cloud Core — comparative structural profile and remedy-prediction register (frozen 22 August 2026). Public digest: snapshot S1, blackholeindex.com/research-snapshots.txt

Savich, I. (2026). Response to the CMA's Consultation on a Proposed Steering Conduct Requirement for Apple's Mobile Platform, 21 July 2026 (published by the CMA 14 August 2026).

Savich, I. (2026). The Black Hole Index: A Structural Lock-In Measurement Framework for Digital Platforms (the 0–10 scoring anchors are published in full in the paper's appendix). SSRN Abstract 6555158, ssrn.com/abstract=6555158.

Black Hole Index (2026). Methodology, scoring rubric and aggregation function. blackholeindex.com/methodology