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 provision: a frame-first structural measurement of switching capability and dependence across four cloud infrastructure providers

Market investigation DMA.100236, Articles 16(1) and 19(1) of Regulation (EU) 2022/1925. One provider-neutral frame frozen on 22 August 2026 before any provider parameter was scored; eight parameters scored, three recorded as not scorable, no composite computed, and no exit-cost quantity produced.

Status: Non-confidential in its entirety — may be published in full. No confidentiality assessment is required before forwarding.

Request: please forward to the team handling DMA.100236.

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

Case: DMA.100236, opened by Commission Decision C(2025) 7948 final of 18 November 2025 under Articles 16(1) and 19(1) of Regulation (EU) 2022/1925

Routing: address given by the EC DMA Team in reply to a procedural enquiry of 1 September 2026

Providers measured: Amazon Web Services · Microsoft Azure · Google Cloud Platform · Oracle Cloud Infrastructure

Record frozen: 22 August 2026 · archive SHA-256 6565778b84b38a9420c41fec37219303548cf6b7fb52f88237c6677b949712fe · content private pending quality assurance, available to the case team on request

Evidence cutoff: 22 August 2026 · first-party provider documentation throughout the scored rows

Independence: no commercial relationship with, and no funding from, any provider measured or any party with an interest in the investigation

1. Introduction, standing and routing

1.1. **This document is non-confidential in its entirety and may be published in full. No confidentiality assessment is required before it is forwarded.** If this is not the team handling market investigation DMA.100236, I ask that it be passed to them.

1.2. I provide this information in connection with DMA.100236, opened by Commission Decision C(2025) 7948 final of 18 November 2025 pursuant to Articles 16(1) and 19(1) of Regulation (EU) 2022/1925, to examine whether the obligations laid down in the Regulation are capable of addressing practices in cloud computing services. It is sent to this address because, in reply to my procedural enquiry of 1 September 2026 asking which channel was appropriate for a non-confidential, unsolicited third-party research information provision referencing DMA.100236, the EC DMA Team replied: “You can use: EC-DMA-EU-CITIZENS@ec.europa.eu”.

1.3. The Black Hole Index is a published measurement instrument for structural dependence on digital platforms: eleven parameters, each scored on published five-anchor rubrics from zero to ten, aggregated by a published function. Its methodology, its scoring rubric and the scores of 184 platforms are public

under CC BY 4.0 at blackholeindex.com. The measurement reported here is a separate exercise, governed by its own frozen boundary, and no earlier published row of that release is an input to it.

1.4. Nothing here is a view on whether any provider should be designated as a gatekeeper, whether any obligation applies to it, or whether any conduct infringes the Regulation. No sentence in this document is a characterisation of any commercial practice. Those are determinations for the Commission; what is offered is a measurement, its method, and the points at which the measurement stopped.

1.5. I am an independent researcher. I have no commercial relationship with, and receive no funding from, Amazon Web Services, Microsoft, Google, Oracle, or any party with an interest in the outcome of this investigation. No party has reviewed, commissioned or been shown this document before its despatch.

1.6. For completeness, my prior contacts with the Commission on adjacent subject matter: a contribution to the targeted consultation on safeguarding the European Union's data sovereignty on 22 August 2026, and a methodological letter to the Joint Research Centre's competence centre on composite indicators on 31 August 2026. Neither is a case submission; both are available on request. Related but separate records have been provided to the United Kingdom's Competition and Markets Authority on 21 July and 31 August 2026 and to the Japan Fair Trade Commission on 30 August 2026. The same record may be relevant to DMA.100032 and DMA.100033, opened the same day as this investigation; I make no submission in those proceedings and take no position on designation.

2. Summary

2.1. Four cloud infrastructure providers — Amazon Web Services, Microsoft Azure, Google Cloud Platform and Oracle Cloud Infrastructure — were measured as four readings of one object under a single provider-neutral frame, frozen on 22 August 2026 before any provider parameter was scored. Eight of the instrument's eleven parameters carry a value for each provider. Three are recorded as *Not Scorable*, with the searches that failed to settle them. No composite score exists for any provider, and none will be published until all eleven are scored.

2.2. The central methodological result is that a rigorous application of this instrument to cloud infrastructure reaches an identification boundary before it reaches a composite. Provider documentation can settle some dimensions of structural dependence; others are properties of the customer–provider relationship and are not settleable from what providers publish; one requires a longitudinal design the frame could not construct without measuring editorial policy instead of capability. That boundary is reported rather than closed by estimation.

2.3. The six provider-attributable findings, in plain words and with the providers named, are these. Configuration state is portable at two providers and proprietary at two, so at Amazon Web Services and Microsoft Azure six years of operational encoding is re-authored on exit rather than carried. Two providers document compatibility with a third's storage protocol in their own materials: Oracle publishes an explicit compatibility interface, Google Cloud documents partial interoperability for "some" tools and libraries, and the protocol is Amazon's. Microsoft Azure alone documents no interoperability at all with that protocol, after a defined search. Amazon Web Services alone documents first-party tooling that writes into a competitor's storage. At Microsoft Azure the management hierarchy must trust the corporate identity tenant and the in-scope build layer requires an organisational construct outside the platform. Three providers document models trained on the customer's own operational history and, in the pages read, document no way to export that learned state.

2.4. One consequence of the frame bears directly on this investigation. Price enters no structural parameter of this instrument by rule. It follows that a remedy which removes or reduces egress charges and changes nothing else moves this structural measure by exactly zero. That is a statement about what this instrument measures, not an argument that such remedies are without value.

3. Method: a frame fixed before any provider parameter was scored

3.1. **Why the frame comes first.** The ordinary way to build a four-provider comparison is to measure the most interesting provider first and then apply the same frame to the others. That method fails in a way the output does not show. Whichever provider is measured first supplies the vocabulary: its catalogue decides which capabilities are in scope, because those are the capabilities in front of the evaluator while the scope is being written. Every later provider is then assessed on how well it matches an architecture drawn, without anyone intending it, around a competitor. This was not hypothetical. A first draft of the boundary named its service surface in one vendor's product names; read as a scope for that vendor it was sound, and read as a template for four it was already contaminated in a way the final table would not have revealed. That draft was withdrawn rather than deleted; it is retained outside the frozen archive and is available on request.

3.2. **What was frozen, and what was amended.** The boundary was frozen on 22 August 2026 before any provider parameter was scored. Two stipulations were found and recorded while scoring — the relational engine and the configuration-authoring posture — and the evidence rules were extended twice, each amendment recorded before the determination it governs. No score changed as a result. The revision histories are part of the record and travel with it.

3.3. **Reference Customer.** One stipulated European enterprise: about 1,200 staff, of whom approximately 150 are engineering and platform; six years' tenure; an estate of approximately 200 terabytes; production under formal change control; a committed-spend agreement; no workforce-productivity suite in scope; a managed open-source relational engine; the provider's first-party configuration surface. Two of these stipulations have declared directions. The open-source engine is favourable to exit, so the portability and substitutability values at that layer are an upper bound: an organisation on a proprietary engine faces a strictly worse position than the one measured here. The exclusion of the workforce-productivity suite narrows the frame in a direction adverse to no provider equally: much of one provider's real-world cross-side pull runs through that excluded estate, and the value scored here is a declared understatement of what a buyer of the full suite experiences.

3.4. **Functional layers.** Twelve layers of the estate, named without vendor vocabulary. Each provider's mapping onto those layers was confirmed cell by cell against first-party documentation — forty-eight cells, all confirmed — before any parameter was scored. A reader who disputes the comparison can therefore dispute a named mapping decision rather than the comparison as a whole.

3.5. **Exit criterion.** A completed exit is defined once, in advance: every workload within the layer inventory runs on a different provider or on the customer's own infrastructure; the data and configuration state required to operate those workloads has been reconstructed there; no business function offered to the customer's own customers has been withdrawn; and the transition completes within twelve months using the Reference Customer's existing engineering staff plus contracted assistance not exceeding twenty-five per cent of that headcount. An export capability is not an exit, and a partial exit is not an exit.

3.6. **Price.** Price enters no structural parameter. The instrument defines the financial exit cost as a separate quantity, to be computed against a dated pricing snapshot — 22 August 2026 for this cohort — and re-

ported alongside the structural profile. **That computation is a queued step. No exit-cost quantity has been produced for this cohort, and this document reports none.**

3.7. **Evidence.** First-party provider documentation throughout the scored rows, read at first hand, quoted with its address and retrieval date, with an evidence cutoff of 22 August 2026. No source dated later may enter any ledger under this boundary. Substitutability is searched symmetrically across the cohort and is cohort-relative by rule. The governing evidence rule is that absence of documentation is not absence of capability: a not-found without a recorded search is not a finding.

3.8. **Neutrality: what the check catches and what it does not.** A script scans the boundary for distinctive vendor service names and fails if it finds any, which catches the visible half of the neutrality claim. It does not catch the harder half: a layer inventory can be shaped by one vendor's architecture without ever naming it. That residual is declared in the boundary rather than asserted away.

4. Results

4.1. Each parameter is scored from zero to ten on published five-anchor rubrics. Six parameters describe capture, where a higher value means greater dependence: data, memory, action, process, network and closeness. Three describe escape, where a higher value means an easier exit: portability, substitutability and human fallback. Two are conditions of the deployment: organisational depth and momentum. The published aggregation is a ratio of a capture product to an escape aggregate; it is not computed here, for the reason given at 4.7. Provenance codes: **P** provider-observed; **R** stipulated by the measurement frame; **R+P** a stipulated baseline, floor or volume, moved or bounded by provider-attributable facts; **PR** relationship-derived and not settleable from provider documentation; **L** longitudinal. A dash is *Not Scorable / Insufficient Evidence*, recorded with its searches.

PARAMETER	SIDE	PROVENANCE	AWS	AZURE	GCP	ORACLE
Data capture (d)	capture	R+P	8	8	8	8
Memory capture (m)	capture	R+P	8	8	7	7
Action capture (a)	capture	R+P	8	8	8	8
Process capture (p)	capture	R	10	10	10	10
Network effects (n)	capture	P	7	5	5	5
Closeness (c)	capture	PR	—	—	—	—
Portability (x)	escape	P	6	5	6	6
Substitutability (s)	escape	P	5	4	5	6
Human fallback (h)	escape	PR	—	—	—	—
Organisational depth (o)	condition	R+P	7	8	7	7
Momentum (t)	condition	L	—	—	—	—

4.2. **What the frame determined, and not the providers.** Process capture reads ten for all four because the boundary places the customer's revenue-bearing product on the platform with no production failover; the top anchor follows deductively from that stipulation, and the parameter carries no discriminating information in this cohort. Organisational depth starts at anchor seven by construction, the

boundary's stipulation of an enterprise-wide, six-year, committed-spend deployment being the anchor's own text; only provider-attributable facts move it. The volume half of data capture and the volume of memory capture are stipulated by the estate and the tenure; the floor of action capture is stipulated by the definition of a managed layer. A table printing these values without their provenance would invite a reader to attribute to providers what the frame decided.

4.3. **What differentiates the providers.** Six findings, each provider-attributable and each evidenced from first-party documentation read within the cutoff.

4.3.1. *Configuration portability.* The first-party configuration services of Google Cloud Platform and Oracle Cloud Infrastructure use a cross-platform format; those of Amazon Web Services and Microsoft Azure are proprietary. In those two cases, and not in the other two, six years of operational encoding is re-authored on exit rather than carried. The caveat travels with the finding: a portable configuration language is not a portable resource graph.

4.3.2. *A storage protocol that others document compatibility with.* Oracle publishes an explicit compatibility interface for Amazon's storage protocol. Google Cloud documents partial interoperability with it, in its own words with "some tools and libraries", which the record treats as the limit of that reading rather than as an equivalent interface. This is first-party evidence, from competitors, that the tool ecosystem orients to that surface. The measurement scored it one anchor step short of the rubric's "de facto industry standard" anchor, and the construct audit records a counter-reading one step lower still and notes it as the thinnest full-step justification among the thirty-two scored rows.

4.3.3. *A missing compatibility surface.* On the reading above, Microsoft Azure alone documents no interoperability with that protocol, after a defined search of its storage documentation set. Applications written against its storage interface are rewritten on exit rather than re-pointed.

4.3.4. *Tooling that writes into a competitor.* Amazon Web Services alone documents first-party tooling that writes into a competitor's storage. The other three document outbound bulk paths to the customer's own infrastructure but no first-party path into a rival platform. The asymmetry lies in that one distinction and is not a claim that the other three document no outbound path.

4.3.5. *Governance bound to the corporate directory.* At Microsoft Azure the management hierarchy must trust the corporate identity tenant — all subscriptions within a management group must trust the same tenant, and the root group's identifier is the tenant identifier — and the in-scope build layer requires an organisational construct outside the platform. Two parameters move on the first of these facts, and the correlation is declared.

4.3.6. *Learned operational state.* Amazon Web Services, Microsoft Azure and Google Cloud Platform each document models trained on the customer's own operational history. In the predictive-autoscaling documentation, read symmetrically for all four providers, no export path for that learned state was found; other classes of derived state were not read for all four and are excluded from this reading. Oracle's autoscaling documentation defines exactly two modes, neither learned.

4.4. **Where equality is the finding.** Action capture reads eight across the cohort: infrastructure-layer autonomy is symmetric, and at the pages read for this determination the estate-level agentic layer was propose-gated at all four providers. **That reading is since qualified.** A first-party page inside the evidence cutoff, which the question set did not reach, documents a customer-selectable mode at one provider in which the agent "investigates and executes actions without waiting for approval"; the documented default at that provider remains approval-gated. The question is pre-registered for all four providers at the

next determination. Data capture reads eight across the cohort because the holdings are stipulated and the observation practices are symmetric at the resolution read.

4.5. **Correlated inputs, named, with their signs.** One documented fact may legitimately drive two parameters, and an aggregation would then weight it twice. Three cases are declared: the configuration-format fact, read once as a path and once as a stock; the storage-compatibility documents; and the identity-plane entanglement. The second of these has opposite signs on different providers: the same documents raise the substitutability of the two providers that document compatibility and evidence the network position of the provider whose protocol it is. A widely implemented interface therefore lowers switching cost at the margin for those who implement it while raising the measured structural position of the party whose interface it is. Whether that is a property of the world or an artefact of this instrument is an open question, stated here rather than resolved.

4.6. **Two scored cells carry declared weaknesses.** Oracle's organisational-depth value rests on a not-found on a getting-started page; a dedicated policy page was not consulted and could raise it by one step. Microsoft Azure's network value is a declared understatement, for the reason given at 3.3. The largest single differential in the table — Amazon Web Services at seven against five for the others on network effects — is therefore qualified on both sides in the record.

4.7. **Why no composite exists.** The protocol forbids computing, displaying or publishing a composite until all eleven parameters are scored. There is no partial composite: the three undetermined parameters enter multiplicatively on both sides of the ratio, two in the capture product and one inside the escape aggregate, so dropping them is an assertion about their values rather than an abstention. Momentum is not scorable because its rubric requires velocity relative to competitors, which requires a common unit over time; the providers' documentation channels exist symmetrically but differ in editorial policy, and an entry count would measure that policy rather than capability. Closeness and human fallback are not scorable because their anchors describe the customer's state — contact frequency, retained capability — which no provider documents. The scoring pipeline enforces this: it first reproduces all published scores of the pinned release, then refuses to compute while any determination is *Not Scorable*.

5. Bearing on the questions in the opening decision

5.1. The opening decision identifies interoperability obstacles, business-user access to data, tying and bundling practices, and excessive fees or other unfair contractual terms. The record bears on these unevenly, and where it bears weakly or not at all I say so.

5.2. **Interoperability.** Sections 4.3.1 to 4.3.4 are interoperability findings of two different kinds, and the distinction may be worth the investigation's attention. Three concern what a provider makes possible at its own interface: a portable configuration format, a compatible storage endpoint, tooling that writes outward. One concerns what competitors do about a surface they do not control. A remedy addressed to interface obligations reaches the first kind. It is not obvious that it reaches the second, and section 4.5 records why that case is also the one where this instrument's own reading is least settled.

5.3. **Access to data.** Section 4.3.6 is the finding with the most bearing on this question, and it concerns a category of data that portability duties are not obviously drafted for. Models trained on a customer's operational history are neither data the customer supplied nor data the customer generated in the ordinary sense; they are state the provider derived from the customer's behaviour. What the record establishes is bounded and is stated as such: in one feature class, read identically for all four providers, no export path was found. That is a not-found after a defined and recorded search, not a demonstration that

no such path exists. Whether such state falls within the reach of an access or portability obligation is a legal question I do not answer.

5.4. **Tying and bundling.** Section 4.3.5 is the finding that bears on this question, and it is narrow. It records two documented architectural dependencies at one provider. It is not a characterisation of any commercial practice, and the record does not offer one.

5.5. **Fees and contractual terms. On this question the record offers no measurement.** Price is excluded from every structural parameter by rule, and the financial exit cost, though defined as a separate quantity with its own snapshot date, has not been computed for this cohort. What the record offers instead is a caution about measurement: a remedy that removes egress charges and changes nothing else moves this structural measure by exactly zero, so progress on a price measure must not be read as progress on a structural one. Which non-price commitments would move which parameters, and in which direction, was pre-registered on 22 August 2026 against the commitments recorded in the United Kingdom Competition and Markets Authority's Board paper *Actions on cloud and business software through the UK digital markets competition regime* of 31 March 2026, before the outcome of that Board's six-month progress review, which was set for late September 2026 and has not been published. The register is specific to that United Kingdom process; I offer it here only as evidence that the predictions were fixed before their test.

5.6. **What the record does not support.** It does not rank the four providers, because no composite exists. It does not establish that any provider holds market power or has abused any position. It does not describe any customer other than the stipulated one; organisations on proprietary relational engines, or on third-party configuration tooling, face different positions, and section 3.3 gives the direction. Applying the instrument's own noise rule to the table, a difference of one anchor step between two providers is not individually quotable; on that rule two differences survive as quotable — Amazon Web Services' network value against the other three, and the Oracle-to-Azure spread on substitutability — and both are qualified at 4.5 and 4.6. The published robustness analysis finds rank ordering stable in aggregate under one-step noise; that is a statement about the ordering, not a licence to quote any individual pairwise gap.

6. Limits of this record, stated by its author

6.1. **Who produced it.** The record's author scored every row and recorded every step of the determination, working under my instruction as principal investigator. Before despatch I have read this document in full and I am responsible for every value and statement in it. Inter-rater reliability for this instrument is not established. A study to establish it has been designed: a cohort of twenty-two systems was drawn by lot from a pinned public release and frozen on 22 August 2026, and the statistical plan — estimators, scale, interval procedure, and thresholds each with a stated consequence — was frozen on 6 September 2026 before any evaluator was approached, with its digest published. The study has not begun. Until it reports, every figure above is one assessor's structured judgement under a published rubric.

6.2. **Anchoring exposure.** The evaluator had previously seen two figures from an earlier release of the instrument: a legacy row labelled AWS at 6.85 and a conglomerate row labelled Microsoft at 6.41. Neither is an input to this record. The exposure is declared because no procedure used here removes it.

6.3. **Earlier company-level rows are a different object.** The public release carries rows labelled AWS and Microsoft produced before this protocol existed, carrying no boundary document. They are not inputs to this measurement and appear in no table here. A figure on the public site under those labels — including a portability value that differs from the one above — describes a different object.

6.4. **Documentation, not behaviour.** Every scored row reads published provider documentation within a declared window. It is a statement about what providers document, not about what any real organisation

experiences.

6.5. **The providers were not contacted.** None of the four was approached, and none has had an opportunity to comment on any value or finding above.

6.6. **Quotations have no retained captures.** The ledgers carry each quotation's address and retrieval date but no page captures. Provider documentation is revised without notice, so a quotation may need to be verified against a web archive rather than a live page. Later determinations retain per-page captures.

6.7. **A material question was not asked.** Whether a provider's identity plane federates outward to third-party services was not in the portability question set for any provider, although first-party material on it falls inside the evidence cutoff. It is pre-registered as a mandatory question of the next determination, put to all four; the predicted movement is zero to one anchor step for providers documenting it.

6.8. **Substitutability is cohort-relative** by rule: no provider's value may be quoted outside the cohort that produced it. **The layer inventory is a judgement call**, recorded in full so that it can be objected to specifically.

6.9. **Evidential weight.** As stated in the boundary, and in the terms previously put to the Commission: no dependency index, this one included, should carry decisive weight in a policy decision on its own.

7. The record itself

7.1. The corpus was archived and frozen on 22 August 2026, before any external review of it and before the outcome of the review referred to at 5.5. It contains fifteen documents and a manifest of per-file digests: the regulatory protocol; the provider-neutral measurement boundary; the evidence rules; the comparability matrix; four provider evidence ledgers; the construct-validity audit of all thirty-two scored rows; the comparative profile; the pre-registered remedy-prediction register; the open research issues; a related boundary note; and two enforcement scripts. Its SHA-256 is 6565778b84b38a9420c41fec37219303548cf6b7fb52f88237c6677b949712fe, published on the day of freezing and independently timestamped the same day.

7.2. **The corpus is not public.** Its content remains private pending completion of quality assurance; only its digest is published, which is what allows anything later released to be checked against the state that existed on the stated date. It is available to the case team on request, in full and at no cost, together with two later items that postdate the archive and are referred to above: the research issue recording the qualification at 4.4, and the register finding recording the unasked question at 6.7.

7.3. Two research notes drawn from the corpus are published in full and may be read without any request, at blackholeindex.com/research/rn-001-portability-is-not-exit-capacity and blackholeindex.com/research/rn-002-cloud-core-identification-boundary. The first sets out why documented portability under-identifies exit capacity; the second why the identification boundary was reported rather than crossed.

7.4. An append-only public record of this work, including every correction and withdrawal made to it, is at blackholeindex.com/research-snapshots.txt. Corrections are appended with their dates; nothing in it is edited after publication.

8. Availability

8.1. I remain available to the case team, at no cost and under no condition, to explain the method, to supply the corpus and the two later items, or to answer a question about any single row of any ledger.

8.2. If any mapping, score, omission or statement above is wrong, I would correct it and record the correction in the public record named at 7.4, as I have done with every earlier correction to this work.

8.3. The one thing I would ask, if it is possible: confirmation that this document has been passed to the team handling DMA.100236.

Ivan Savich

CREATOR, BLACK HOLE INDEX · INDEPENDENT RESEARCHER

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

References

European Commission (2025). Commission Decision C(2025) 7948 final of 18 November 2025 opening a market investigation pursuant to Articles 16(1) and 19(1) of Regulation (EU) 2022/1925, Case DMA.100236.

Regulation (EU) 2022/1925 of the European Parliament and of the Council of 14 September 2022 on contestable and fair markets in the digital sector (Digital Markets Act), Articles 16 and 19.

Black Hole Index (2026). Cloud Core — provider-neutral measurement boundary v1.3, evidence rules, comparability matrix and four provider evidence ledgers (frozen 22 August 2026; archive SHA-256 6565778b84b38a9420c41fec37219303548cf6b7fb52f88237c6677b949712fe).

Black Hole Index (2026). Cloud Core — comparative structural profile and pre-registered remedy-prediction register (frozen 22 August 2026).

Savich, I. (2026). Portability is not exit capacity. Research note RN-001. blackholeindex.com/research/rn-001-portability-is-not-exit-capacity

Savich, I. (2026). Why Cloud Core refused a composite score. Research note RN-002. blackholeindex.com/research/rn-002-cloud-core-identification-boundary

Savich, I. (2026). The Black Hole Index: A Structural Lock-In Measurement Framework for Digital Platforms (scoring anchors and aggregation). blackholeindex.com/paper

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

Black Hole Index (2026). Append-only public research record. blackholeindex.com/research-snapshots.txt