Methodology Working Group
Category: Standards Track
Author: MarinaDNS (R. Mastiyev, Editor)

Standard ID: NODA-v1.0.1
Status: First Release, Revision 1
Published: August 2026
Revised: October 2026
Obsoletes: NODA-v1.0
Updates: none

NETWORK OPERATIONAL DOMAIN ASSESSMENT (NODA)

Methodology Specification — Version 1.0.1

Status of This Document. This is the first public release of this methodology. Scores, reports, and findings produced under NODA-v1.0 are citable as of the date of publication above. Distribution of this document is unrestricted. All provisions herein are subject to the versioning discipline defined in Section 12. Version 1.0.1 is an editorial revision under Section 12 with zero scoring impact; every change and its reason is recorded in Appendix C.

1. Abstract

The Network Operational Domain Assessment (NODA) defines a reproducible, quantitative method for evaluating the operational quality, security posture, and technical correctness of a domain's publicly observable Domain Name System (DNS) infrastructure. NODA is issued by MarinaDNS as an open, versioned methodology, intended for use by domain operators, researchers, and automated tooling.

NODA produces two independent composite scores per assessment — a Health score and a Best Practice score — rather than a single blended figure. The Health axis measures whether the assessed infrastructure functions correctly and presents no active operational risk. The Best Practice axis measures conformance to recommended standards and hardening practices that reduce future or latent risk. A domain MAY score highly on one axis while scoring poorly on the other; this specification treats that outcome as informative, not contradictory.

This document defines terminology, conformance requirements, assessment scope, the assessment procedure, the scoring formula, category and severity taxonomies, report structure, and the versioning discipline that governs all future revisions of this methodology.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119.

— Check — a single, deterministic test executed against a domain's public DNS infrastructure, producing exactly one Finding.

— Finding — the output of a Check: a status (Pass, Warning, Critical, Informational, or Not Applicable), together with a human-readable explanation and, where applicable, remediation guidance.

— Category — a grouping of Checks by subject-matter domain (Section 7). Every Check belongs to exactly one Category.

— Axis — a grouping of Checks by the type of quality they measure: Health (current function) or Best Practice (standards conformance). A Check MAY contribute to one axis or be apportioned across both (Section 6).

— Applicability — a fixed property of a Check indicating whether it is Unconditional (always evaluated) or Conditional (evaluated only when a prerequisite record set is present in the target's configuration). A Conditional Check resolves to a Not Applicable Finding when its prerequisite is absent (Section 6.3).

— Composite Score — the weighted aggregate of Finding compliance values within a single axis, expressed as an integer from 0 to 100 (Section 6).

— Report — the complete, immutable output of one assessment, comprising all Findings, both Composite Scores, and required metadata (Section 9).

3. Conformance

An implementation of this methodology is conformant if it satisfies every MUST and REQUIRED provision in this document. Conformance SHOULD be re-verified whenever the implementation is modified. Any system that publishes or displays a NODA score MUST disclose the methodology version (Section 12) under which that score was produced.

This document does not itself constitute a Report. A Report is conformant only if it was generated by a conformant implementation.

4. Scope

The assessment scope of NODA includes: authoritative DNS delegation and glue, zone configuration (SOA, A, MX), email authentication records (SPF, DKIM, DMARC), DNS security extensions (DNSSEC), and operational consistency and resilience signals derivable from public DNS and mail-exchange data.

The following are explicitly out of scope, and their absence from a Report MUST NOT be interpreted as a compliance signal in either direction: internal or split-horizon DNS, registrar account security, application-layer security of hosted services, distributed denial-of-service mitigation, and physical infrastructure security. Assessment MUST NOT extend to non-DNS-observable properties of the target; this boundary exists to prevent lateral assessment drift.

5. Assessment Methodology

This section is procedural. It describes how data is collected and MUST be read independently of Section 6, which defines what the collected data means.

5.1 Design Principles

— Objective — every Check is deterministic; identical infrastructure MUST produce identical Findings.

— Transparent — every Finding that is not a Pass MUST be accompanied by an explanation sufficient for the operator to reproduce the deduction independently. NODA reserves no proprietary or undisclosed scoring logic.

— Actionable — every non-Pass Finding SHOULD include remediation guidance (Section 9).

— Vendor-neutral — Checks are derived from published, open standards, not from any single vendor's implementation.

— Versioned — the methodology evolves under the discipline of Section 12, without invalidating prior Reports.

5.2 Procedure

Assessment of a domain proceeds through three phases, executed in order:

— Discovery — passive resolution of the domain's authoritative delegation chain via a recursive resolver operated by the implementation.

— Enumeration — collection of the record set relevant to each in-scope Category: NS, SOA, A, MX, TXT (SPF/DMARC), DNSKEY/DS (DNSSEC), and associated mail-exchange authentication records.

— Verification — execution of each Check in Appendix A against the enumerated data, producing one Finding per Check.

Implementations SHOULD query authoritative and recursive infrastructure operated or controlled by MarinaDNS rather than third-party public resolvers, to avoid attributing query load to the assessed domain's infrastructure or violating third-party operator policy.

NODA checks may be initiated through noda.marinadns.io directly, or through other MarinaDNS-family front-ends that query the NODA API on the operator's behalf. Regardless of origin, all checks run against the same scoring engine and are logged to a single, canonical dataset — results and methodology are identical across every surface.

6. Scoring Model

NODA computes two independent Composite Scores per Report: a Health score (H) and a Best Practice score (B). The two scores MUST be reported independently and MUST NOT be blended into a single figure at any stage.

Each Check i is assigned four fixed properties, published in Appendix A: a Category weight w_i (Section 7), an axis allocation pair (h_i, b_i) where $h_i + b_i = 1$ — except for Checks designated Informational-Only (Appendix A.6), which carry no axis allocation and are permanently excluded from both Composite Score formulas — an Applicability designation (Section 6.3), and a compliance function that maps its Finding status to a value $V_i \in \{0, 0.5, 1\}$: Critical = 0, Warning = 0.5, Pass = 1. Findings of status Informational or Not Applicable are excluded from scoring entirely; the two statuses are excluded for distinct reasons and MUST NOT be conflated in a Report's presentation layer — Informational denotes a Finding that was evaluated and found to carry no scoring weight, while Not Applicable denotes a Check that was not evaluated at all because its prerequisite condition was not met (Section 6.3).

For the majority of Checks, the axis allocation is (1, 0) — Health only — or (0, 1) — Best Practice only. A small number of Checks that measure a refinement of a redundancy signal already captured elsewhere by a baseline existence Check are apportioned across both axes (Section 6.2).

6.1 Composite Formula

$$H = \Sigma (w_i \cdot h_i \cdot V_i) \ / \ \Sigma (w_i \cdot h_i) \quad [\text{scored over all } i \text{ with } h_i > 0]$$

$$B = \Sigma (w_i \cdot b_i \cdot V_i) \ / \ \Sigma (w_i \cdot b_i) \quad [\text{scored over all } i \text{ with } b_i > 0]$$

Both H and B are computed on a 0–1 scale and then expressed as integers from 0 to 100 by multiplying by 100 and rounding to the nearest integer, with halves rounded away from zero. An axis whose denominator is zero has no Composite Score and is reported as Unscored. A Check with a split allocation contributes to both sums simultaneously, weighted by its allocation fraction; this prevents a single underlying signal from being counted twice at full weight on either axis.

6.2 Split-Allocation Rule

A Check that refines a redundancy or diversity signal already captured, at the level of mere existence, by a separate baseline Check MUST be apportioned 30% to the Health axis and 70% to the Best Practice axis. This rule exists to avoid double-counting: the baseline existence of redundancy is already scored in full on the Health axis by another Check, so the refinement Check's contribution to Health is deliberately minority-weighted, while its primary contribution — measurement of standards-conformant redundancy depth — is scored on the Best Practice axis. The 30% Health allocation reflects that concentration of redundant components within a single subnet or autonomous system is itself a correlated-failure exposure at the time of assessment, independent of recommended depth. NS Different Subnets and ASN Diversity, whose baseline existence Check is Nameserver Count (CAT-INFRA), measure distinct failure domains — routing prefix and autonomous system respectively — and are scored independently.

A Check that measures whether a redundancy claim is factually true at the time of assessment, rather than whether it conforms to a recommended depth, is scored on the Health axis only, regardless of its relationship to a baseline existence Check. The distinction is whether the Check answers “does this work today” (Health) or “does this meet a recommended standard beyond bare function” (Best Practice).

6.3 Conditional Applicability

Not every Check applies to every domain by design, not by omission. A domain that operates no mail service has no MX record set to evaluate; this is a legitimate configuration choice under RFC 5321 and RFC 1035, not a deficiency, and MUST NOT be scored as one.

Each Check carries a fixed Applicability designation of Unconditional or Conditional. An Unconditional Check is evaluated for every assessed domain regardless of configuration. A Conditional Check is evaluated only when its declared prerequisite record set is present; when the prerequisite is absent, the Check MUST resolve to a Not Applicable Finding rather than a Critical or Warning Finding, and is excluded from both the numerator and denominator of whichever axis formula (Section 6.1) it would otherwise have contributed to.

This exclusion does not alter the weight w_i carried by any other Check: the remaining Unconditional and applicable Conditional Checks within that Category continue to carry the Category's published weight w_i individually, and no weight is redistributed to other Categories or Checks. Because w_i applies per Check (Section 7), the Category's effective share of the affected axis decreases by the weight of the excluded Checks; this is intended, since a Check that was not evaluated contributes no evidence in either direction.

Within CAT-MAIL and CAT-RELI, the prerequisite for all mail-exchange-dependent Conditional Checks is the presence of at least one MX Record. The MX Records Check itself is Unconditional and MUST treat the deliberate absence of any MX record as a valid Pass state, not a failure; its function is to validate the MX configuration if one is declared, not to require that one exist. SPF Record and DMARC Record remain Unconditional irrespective of MX presence: RFC 7208 §10.1 recommends a null SPF policy for domains that send no mail, and an equivalent DMARC policy serves the same anti-spoofing purpose, making both checks independently meaningful on a domain with no mail service at all. DKIM Record is Conditional, since no signing key is meaningful in the absence of an outbound mail path. DKIM selector names are not standardized by RFC 6376 and cannot be reliably enumerated through passive DNS lookup alone; where the DKIM Record Check is applicable but no key is found among the selectors checked, it resolves to an Informational Finding rather than a Critical Finding, since absence of a signature at a guessed selector is not equivalent to a confirmed absence of DKIM signing. The selectors checked are listed in the Finding's explanation in each Report. Because non-detection resolves to Informational, the DKIM Record Check can raise but never lower a Composite Score: two domains with equivalent DKIM deployment may differ in Health score according to whether their selector is among those checked.

Within CAT-SEC, the prerequisite for DNSSEC-dependent Conditional Checks is the presence of at least one DNSKEY record in the assessed zone. The DNSSEC Check itself is Unconditional and evaluates whether a zone is signed at all; Checks that examine the properties of an existing signature — DNSSEC Multi-Algorithm Signing and DNSSEC Algorithm Ban List — are Conditional on that signature already existing, and resolve to a Not Applicable Finding when the zone is unsigned.

The Blacklist Check (CAT-OPS) is Unconditional but evaluates mail-exchange IP addresses only. When no MX Record is present there is no address to evaluate, and the Check resolves to an Informational Finding; CAT-OPS then contributes nothing to the Health axis for that domain.

6.4 Interpretation Tiers

Each Composite Score is accompanied in a Report by an interpretation tier, assigned independently per axis: 90–100 Exceptional; 75–89 Good; 60–74 Needs Improvement; 40–59 Poor; 0–39 Critical. An axis with no Composite Score (Section 6.1) is labelled Unscored. Tiers are a presentation aid derived solely from the integer Composite Score; they carry no scoring weight and do not alter either score. The tier label Critical is distinct from the Critical Finding status (Section 6) and from severity level S1 (Section 8).

7. Category Definitions

Category and Axis are independent dimensions. Category groups Checks by subject matter and carries the fixed weight w_i used in Section 6. Axis groups Checks by the type of quality measured, per Check, regardless of Category. The Category named “Operational Best Practices” (CAT-OPS) is a subject-matter grouping and MUST NOT be conflated with the Best Practice axis; a Check in any Category MAY be allocated to the Best Practice axis, and CAT-OPS Checks MAY themselves be allocated to the Health axis.

Category ID	Definition	Weight
CAT-INFRA	DNS Infrastructure — authoritative delegation, glue, SOA, and address record correctness	30%
CAT-SEC	Security — DNSSEC and related cryptographic protections	25%

Category ID	Definition	Weight
CAT-MAIL	Email Authentication — SPF, DKIM, DMARC presence and configuration	20%
CAT-RELI	Reliability — mail-exchange and nameserver redundancy, diversity, and consistency	15%
CAT-OPS	Operational Best Practices — reputation and operational hygiene signals	10%

Category weights sum to 100% and apply identically within both the Health and Best Practice formulas (Section 6.1). Each weight w_i is a per-Check multiplier: every Check in a Category carries that Category's weight individually in the Composite Formula. A Category's effective share of an axis is therefore the sum of $w_i \cdot h_i$ (or $w_i \cdot b_i$) over its applicable Checks divided by that axis's denominator, and depends both on the number of Checks the Category contributes to the axis and on which Conditional Checks apply to the assessed domain. The published weights express the relative importance of an individual Check in each Category, not a fixed share of either Composite Score; Appendix D tabulates effective shares for two reference configurations. A Category with no Checks currently allocated to a given axis contributes zero to that axis's denominator and is excluded from that axis's computation, not scored as zero. CAT-OPS currently carries a single Check; additional CAT-OPS Checks are candidates for future MINOR revisions (Section 13).

8. Severity Classification

Every non-Pass Finding MUST be assigned exactly one of the following four severity levels, independent of its axis or Category:

Level	Label	Definition
S1	Critical	The Check's failure condition is met in full. On a Health-axis Check this indicates likely active service impact or exposure, and remediation SHOULD be treated as urgent; on a Best-Practice-axis Check it indicates the complete absence of a recommended protection.
S2	Warning	An avoidable operational risk or a departure from recommended practice.
S3	Recommendation	A best-practice improvement opportunity with no current operational impact.
S4	Informational	Contextual observation only. Excluded from scoring (Section 6).

A Not Applicable Finding (Section 6.3) carries no severity classification and, like Informational, is excluded from scoring. It differs from Informational in that no evaluation of the underlying condition occurred; Informational Findings reflect an observation, Not Applicable Findings reflect a Check that did not run.

Status and severity are distinct properties. Status (Pass, Warning, Critical, Informational, Not Applicable) alone determines the compliance value V_i (Section 6); severity is a presentation label and has no effect on either Composite Score. A Critical status always carries S1 and an Informational status carries S4. A Warning status carries S3 for the following Checks and S2 for all others: Parent NS Records, Nameserver Count, SOA MNAME, SOA Serial Format, WWW Record, DNSSEC, and DMARC Record.

9. Report Generation

A Report MUST be packaged as a single self-contained JSON object and MUST include: a Report Identifier, the methodology version under which it was produced, the full set of Findings with their Category, Axis, and severity, and both Composite Scores computed independently.

The Report Identifier MUST be composed of the assessed domain, the ISO 8601 UTC timestamp of assessment, and the methodology version, and MUST NOT be reused. Every non-Pass Finding included in a Report SHOULD carry remediation guidance consisting of a technical explanation, an operational-impact statement, and a step-by-step remediation procedure with references to the relevant standards (Section 14).

Per-Check short-form identifiers are RESERVED for introduction in a future MINOR revision (Section 12) and MAY be added to Report output at that time without altering the meaning or values of any field defined in this version.

10. RFC Adherence Ledger

Independent of the Category and Axis taxonomies (Sections 6–7), every Check in Appendix A is annotated with the specific RFC provisions it verifies, restricted to citations using the normative key words MUST or SHALL (Section 2). This Ledger exists to make each Check's standards basis independently auditable without requiring the reader to trace it through remediation prose.

A citation using SHOULD, RECOMMENDED, or MAY is never listed in the Ledger; such provisions are reflected in remediation guidance (Section 9). The Ledger records normative basis only and does not determine a Finding's status or severity, which are set by each Check's compliance function (Section 6) and the definitions of Section 8; a Check MAY produce a Critical or Warning Finding without a MUST/SHALL citation, and such Checks are recorded in the Ledger as having an operational-practice basis. Where a cited document is advisory in nature rather than a standards-track RFC — as with RFC 5358's guidance on reflector-attack prevention, relevant to the Open Recursor Check — the Ledger records it as advisory and does not treat it as a MUST/SHALL source.

The Ledger is descriptive of the Appendix A Check set and carries no independent scoring weight — it MUST NOT be blended into either Composite Score (Section 6).

11. Historical Consistency

Reports are permanent once issued. A Report MUST render, indefinitely, under the scoring rules of the methodology version recorded in that Report's metadata. A later methodology version MUST NOT be applied retroactively to an existing Report under any circumstance; a domain MAY be reassessed to obtain a new Report under the current version, but the prior Report MUST remain accessible and unaltered.

12. Versioning Policy

This methodology follows Semantic Versioning (SemVer 2.0.0).

— MAJOR — a change to the axis model itself, the Category set, the Composite Formula (Section 6.1), any change to a published property of an existing Check (Category, weight, Axis allocation, Applicability, or compliance function), or a Category weight adjustment. Following the first public Report issued under this methodology, a MAJOR revision MUST NOT be issued more than once per twelve-month period, in order to preserve the comparability of previously published scores.

— MINOR — introduction of a new Check with all of its published properties assigned at introduction. MINOR revisions MUST NOT alter any published property of a Check that existed under the prior version. Because a new Check enters the denominator of each axis to which it is allocated, Composite Scores are directly comparable only between Reports produced under the same MAJOR.MINOR version.

— PATCH — editorial clarification, or a change to Report presentation labels, with zero impact on any Composite Score. Reports produced under different PATCH versions of the same MAJOR.MINOR version are score-equivalent.

This constraint is self-imposed and stated here in the same normative terms as the rest of this document; a maintainer bound by a published limit on how often it may revise foundational structure is itself a signal of methodological seriousness, not a limitation of one.

13. Limitations

NODA assesses only publicly observable DNS infrastructure at the moment of assessment. It does not constitute a complete security audit, does not evaluate physical infrastructure or social-engineering vectors, and cannot guarantee protection against vulnerabilities not yet publicly known.

During pre-release drafting, before the first Report was issued under this methodology, six Checks were added to the original catalogue: Open Recursor and AXFR / Zone Transfer (CAT-SEC), ASN Diversity (CAT-RELI), SPF Lookup Limit (CAT-MAIL), and DNSSEC Multi-Algorithm Signing and DNSSEC Algorithm Ban List (CAT-SEC), all catalogued in Appendix A. These additions preceded first release and do not constitute MINOR revisions under Section 12. CAA record presence, IPv6/AAAA availability, and DNS TTL variance remain candidates for future MINOR revisions and have not yet been implemented. At the time of this version's publication, CAT-OPS remains under-populated relative to the other four Categories (Section 7) and is the priority for further expansion.

This version scores the presence of a DNSSEC signature (DNSSEC) and the acceptability of its algorithm (DNSSEC Algorithm Ban List) as separate Checks and does not cap one by the other. A zone signed exclusively with an algorithm prohibited for signing by RFC 9904 may therefore receive a higher Best Practice score than an unsigned zone.

Category and Axis assignment reflect an editorial judgment made by the maintainer, published in full in Appendix A, and subject to revision only under the discipline of Section 12.

14. References

[RFC1034] Mockapetris, P., "Domain Names — Concepts and Facilities," November 1987.

[RFC1035] Mockapetris, P., "Domain Names — Implementation and Specification," November 1987.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," March 1997.

[RFC2181] Elz, R. and Bush, R., "Clarifications to the DNS Specification," July 1997.

[RFC2182] Elz, R. et al., "Selection and Operation of Secondary DNS Servers," July 1997.

[RFC2308] Andrews, M., "Negative Caching of DNS Queries (DNS NCACHE)," March 1998.

[RFC4033] Arends, R. et al., "DNS Security Introduction and Requirements," March 2005.

[RFC4034] Arends, R. et al., "Resource Records for the DNS Security Extensions," March 2005.

[RFC4035] Arends, R. et al., "Protocol Modifications for the DNS Security Extensions," March 2005.

[RFC5321] Klensin, J., "Simple Mail Transfer Protocol," October 2008.

[RFC5358] Damas, J. and Neves, F., "Preventing Use of Recursive Nameservers in Reflector Attacks," October 2008. Advisory; not a MUST/SHALL source (Section 10).

[RFC5936] Lewis, E. and Hoenes, A., "DNS Zone Transfer Protocol (AXFR)," June 2010.

[RFC6376] Crocker, D. et al., "DomainKeys Identified Mail (DKIM) Signatures," September 2011.

[RFC6781] Kolkman, O., Mekking, W., and Gieben, R., "DNSSEC Operational Practices, Version 2," December 2012. Advisory; not a MUST/SHALL source (Section 10). Cited in Appendix A remediation guidance for DNSSEC Multi-Algorithm Signing regarding the disputed scope of multi-signature requirements across zone RRsets.

[RFC7208] Kitterman, S., "Sender Policy Framework (SPF)," April 2014.

[RFC7766] Dickinson, J., Dickinson, S., Bellis, R., Mankin, A., and Wessels, D., "DNS Transport over TCP - Implementation Requirements," March 2016.

[RFC9471] Andrews, M., Huque, S., Wouters, P., and Wessels, D., "DNS Glue Requirements in Referral Responses," September 2023.

[RFC9904] Hardaker, W. and Kumari, W., "DNSSEC Cryptographic Algorithm Recommendation Update Process," November 2025. (Obsoletes RFC 8624.)

[RFC9991] Kucherawy, M. et al., "Domain-based Message Authentication, Reporting, and Conformance (DMARC)," March 2025. (Obsoletes RFC 7489.)

A. Appendix A — Check Catalogue

The following table enumerates every Check defined under this version. Axis allocation is expressed as Health/Best Practice. Applicability of “Conditional (MX)” denotes a Check that resolves to Not Applicable when the target domain has no MX Records (Section 6.3); “Conditional (DNSSEC)” denotes a Check that resolves to Not Applicable when the target domain’s zone has no DNSKEY record. Short-form Check identifiers are reserved (Section 9) and not yet assigned; Checks are referenced by title pending that MINOR revision. Appendix A.6 separately catalogues Informational-Only Checks, which carry no axis allocation and are never scored. The RFC basis column is the RFC Adherence Ledger of Section 10. It cites the specific provision a Check verifies where that provision is a MUST, SHALL or MUST NOT requirement; "normative, pre-RFC 2119" and "normative, no RFC 2119 keyword" mark binding requirements stated in plain wording in documents that predate or deliberately avoid RFC 2119 key words; "Advisory" marks a non-binding source; "Operational practice" marks a Check with no MUST/SHALL basis (Section 10).

A.1 CAT-INFRA — DNS Infrastructure (30%)

Check	Axis (H/B)	Applicability	RFC basis
Parent NS Records	100/0	Unconditional	Operational practice
Nameservers Authoritative	100/0	Unconditional	Operational practice
DNS Parent Sent Glue	100/0	Unconditional	RFC 9471 §3.1 (MUST)
Same Glue	100/0	Unconditional	Operational practice
NS Consistency	100/0	Unconditional	Operational practice
Nameserver Count	100/0	Unconditional	RFC 1034 §4.1 (normative, pre-RFC 2119)
NS TCP	100/0	Unconditional	RFC 7766 §5 (MUST)
NS CNAME Check	100/0	Unconditional	RFC 2181 §10.3 (normative, no RFC 2119 keyword)
NS IPs Public	100/0	Unconditional	Operational practice
SOA Record	100/0	Unconditional	Operational practice
SOA MNAME	100/0	Unconditional	Operational practice
SOA Serial Format	0/100	Unconditional	Operational practice
SOA Serial Sync	100/0	Unconditional	Operational practice
A Record	100/0	Unconditional	Operational practice
A Record IPs Public	100/0	Unconditional	Operational practice
WWW Record	0/100	Unconditional	Operational practice

A.2 CAT-SEC — Security (25%)

Check	Axis (H/B)	Applicability	RFC basis
DNSSEC	0/100	Unconditional	Operational practice
Open Recursor	100/0	Unconditional	Advisory: RFC 5358 (BCP 140)
AXFR / Zone Transfer	100/0	Unconditional	Operational practice
DNSSEC Multi-Algorithm Signing	0/100	Conditional (DNSSEC)	RFC 4035 §2.2 (MUST)
DNSSEC Algorithm Ban List	0/100	Conditional (DNSSEC)	RFC 9904 §3, Table 2 (MUST NOT)

DNSSEC Multi-Algorithm Signing passes when every algorithm present in the zone's apex DNSKEY RRset has at least one covering RRSIG over that RRset (RFC 4035 §2.2); a zone signed with a single algorithm therefore passes.

A.3 CAT-MAIL — Email Authentication (20%)

Check	Axis (H/B)	Applicability	RFC basis
SPF Record	100/0	Unconditional — scored even with no MX	RFC 7208 §3.2 (MUST NOT)
SPF Lookup Limit	100/0	Unconditional	RFC 7208 §4.6.4 (MUST)
DKIM Record	100/0	Conditional (MX)	Operational practice
DMARC Record	100/0	Unconditional — scored even with no MX	Operational practice

A.4 CAT-RELI — Reliability (15%)

Check	Axis (H/B)	Applicability	RFC basis
NS Different Subnets	30/70	Unconditional	Operational practice
ASN Diversity	30/70	Unconditional	Operational practice
MX Records	100/0	Unconditional — absence of MX is a valid Pass	Operational practice
MX Validity	100/0	Conditional (MX)	RFC 5321 §2.3.5 (MUST NOT)
MX Not IP	100/0	Conditional (MX)	RFC 1035 §3.3.9 (normative, pre-RFC 2119)
MX CNAME	100/0	Conditional (MX)	RFC 2181 §10.3 (normative, no RFC 2119 keyword)
MX Consistency	100/0	Conditional (MX)	Operational practice
MX Redundancy	30/70	Conditional (MX)	Operational practice
Duplicate MX IPs	100/0	Conditional (MX)	Operational practice
MX IPs Public	100/0	Conditional (MX)	Operational practice
MX Reverse DNS (PTR)	0/100	Conditional (MX)	Operational practice

A.5 CAT-OPS — Operational Best Practices (10%)

Check	Axis (H/B)	Applicability	RFC basis
Blacklist Check	100/0	Unconditional	Operational practice

A.6 Informational-Only Checks (Non-Scored)

Checks in this section are never scored and carry no axis allocation (Section 6). Their Finding status is fixed at Informational by design, regardless of the underlying configuration value observed — this distinguishes them from scored Checks that can sometimes produce an Informational Finding at runtime (e.g. Blacklist Check when no MX record exists to evaluate). Each entry below carries an informal origin-category tag for grouping purposes only; this tag has no bearing on scoring, since Category weight (Section 7) does not apply to Informational-Only Checks. New Checks MAY be added to this section as a PATCH revision (Section 12), since doing so has zero scoring impact.

Check	Axis (H/B)	Applicability	Origin Tag	RFC basis
SOA Refresh	—/—	Unconditional	Infra-adjacent	Operational practice
SOA Retry	—/—	Unconditional	Infra-adjacent	Operational practice
SOA Expire	—/—	Unconditional	Infra-adjacent	Operational practice
SOA Minimum TTL	—/—	Unconditional	Infra-adjacent	Operational practice

Total Checks defined under this version: 41 (37 scored, 4 Informational-Only per Appendix A.6). Of the 37 scored Checks, 9 are Conditional on the presence of at least one MX Record (8 within CAT-RELI, 1 within CAT-MAIL), 2 are Conditional on the presence of a DNSKEY record (both within CAT-SEC), and the remaining 26 are Unconditional. The 4 Informational-Only Checks are all Unconditional.

B. Appendix B — Scoring Examples

The following worked example illustrates the Composite Formula (Section 6.1) over a representative subset of eleven Checks for a hypothetical domain, example.org. It is illustrative only; a production Report is computed over the full Check set defined in Appendix A.

Check	Category (w)	Axis (h/b)	Finding	V_i
Parent NS Records	INFRA (.30)	1.0 / 0.0	Pass	1.0
SOA Serial Sync	INFRA (.30)	1.0 / 0.0	Pass	1.0
SOA Serial Format	INFRA (.30)	0.0 / 1.0	Warning	0.5
WWW Record	INFRA (.30)	0.0 / 1.0	Pass	1.0
DNSSEC	SEC (.25)	0.0 / 1.0	Critical	0.0
SPF Record	MAIL (.20)	1.0 / 0.0	Pass	1.0
DMARC Record	MAIL (.20)	1.0 / 0.0	Pass	1.0
MX Redundancy	RELI (.15)	0.3 / 0.7	Warning	0.5
Duplicate MX IPs	RELI (.15)	1.0 / 0.0	Pass	1.0
NS Different Subnets	RELI (.15)	0.3 / 0.7	Critical	0.0
Blacklist Check	OPS (.10)	1.0 / 0.0	Pass	1.0

Applying Section 6.1 across these eleven Checks:

H = 1.2725 / 1.3400 = 0.9496 → Health Score: 95

B = 0.5025 / 1.0600 = 0.4740 → Best Practice Score: 47

Per the interpretation tiers in Section 6.4, a Health Score of 95 reads as Exceptional — the domain's DNS infrastructure functions correctly with no active exposure among the Checks sampled. A Best Practice Score of 47 reads as Poor — principally driven by the absence of DNSSEC and the non-diversified nameserver placement. This is the intended output of the dual-axis model: the domain is not at risk today, but has clear, itemized headroom against recommended standards, and the Report communicates both facts without collapsing them into one number that would obscure either.

C. Appendix C — Revision History

This appendix records every revision of this methodology. Each entry states the revision type under Section 12 and its scoring impact.

Version	Type	Summary
1.0 (August 2026)	First Release	Initial publication.
1.0.1 (October 2026)	PATCH	Editorial corrections with zero scoring impact; see C.1.

C.1 Changes in Version 1.0.1

— Section 6.1 — Stated the rounding rule used to express Composite Scores as integers, and the Unscored case. Reason: the rule was applied but not documented, which impaired independent reproduction (Section 5.1).

— Section 6.2 — Removed the requirement that the baseline Check be in the same Category, and stated the basis for the 30% Health allocation. Reason: NS Different Subnets and ASN Diversity, whose baseline Check is in CAT-INFRA, did not meet the rule as written.

— Section 6.3 — Corrected the statement that excluding a Conditional Check does not reduce its Category's weight. Reason: under Section 6.1 weights apply per Check, so exclusion reduces the Category's effective share; the formula is unchanged.

— Section 6.3 — Disclosed that the DKIM Record Check can raise but not lower a score, and the Blacklist Check's behaviour when no MX Record exists. Reason: both behaviours were in effect but undocumented.

— Section 6.4 (new) — Published the interpretation tiers used in Reports. Reason: tiers were applied in Reports and referenced in Appendix B but not defined.

— Section 7 — Clarified that Category weights are per-Check multipliers, added Appendix D, and removed text describing the pre-release state. Reason: the previous wording implied fixed Category shares that the Composite Formula does not produce.

— Section 8 — Extended the S1 definition to Best-Practice-axis Checks and stated the mapping between status and severity. Reason: the S1 definition did not describe Critical Findings on Best-Practice-axis Checks, and the mapping was undocumented.

— Section 10 — Clarified that the Ledger records normative basis and does not determine status or severity. Reason: the previous wording, read literally, excluded Critical or Warning Findings for Checks without a MUST/SHALL citation, contrary to the Check catalogue.

— Section 12 — Moved Category weight adjustment from MINOR to MAJOR, defined the MINOR constraint in terms of published Check properties, and stated report comparability across versions. Reason: the previous MINOR definition contradicted itself; the policy is made stricter, not looser.

— Section 13 — Recorded that six Checks were added before first release, and disclosed the separate scoring of DNSSEC signature presence and algorithm acceptability. Reason: the previous wording implied MINOR revisions had occurred under version 1.0.

— Appendix A — Stated the pass condition of DNSSEC Multi-Algorithm Signing. Reason: the Check title could be read as requiring more than one algorithm.

— Appendix B — Corrected the Best Practice tier of the worked example from Needs Improvement to Poor, and referenced Section 6.4. Reason: alignment with the tiers applied in Reports.

— Appendix A — Added the RFC basis column, which publishes the RFC Adherence Ledger for every Check. Reason: Section 10 stated that every Check in Appendix A was annotated with the provisions it verifies, but no such annotation was published.

— Sections 13 and 14 — Replaced RFC 8624 with RFC 9904, which obsoleted it in November 2025, and added RFC 1034, RFC 7766 and RFC 9471 as Ledger sources. Reason: RFC 8624 was already obsolete at first release; the new references support the Ledger entries.

— Header, title, and Status of This Document — Version updated to 1.0.1.

None of these changes alters any Check property, any compliance value, or any Composite Score. Reports issued under version 1.0 remain valid and score-equivalent (Sections 11 and 12).

D. Appendix D — Effective Category Shares (Informative)

This appendix is informative. For two reference configurations it shows each Category's effective share of each axis: the sum of $w_i \cdot h_i$ (or $w_i \cdot b_i$) over the Category's applicable Checks, divided by the axis denominator (Section 7). Shares are rounded to one decimal place. They describe the Composite Formula and do not modify it.

D.1 MX Records present, zone signed

Category	Effective share of Health	Effective share of Best Practice
CAT-INFRA (30%)	61.9%	33.1%
CAT-SEC (25%)	7.4%	41.3%
CAT-MAIL (20%)	11.8%	— (no Checks allocated)
CAT-RELI (15%)	17.5%	25.6%
CAT-OPS (10%)	1.5%	— (no Checks allocated)

D.2 No MX Records, zone unsigned

Category	Effective share of Health	Effective share of Best Practice
CAT-INFRA (30%)	75.8%	56.6%
CAT-SEC (25%)	9.0%	23.6%
CAT-MAIL (20%)	10.8%	— (no Checks allocated)
CAT-RELI (15%)	4.3%	19.8%
CAT-OPS (10%)	— (Blacklist Check is Informational)	— (no Checks allocated)