

Methodology Working Group
Category: Standards Track
Author: MarinaDNS (R. Mastiyev, Editor)
Obsoletes: none

Standard ID: NODA-v1.0
Status: First Release
Published: August 2026
Updates: none

NETWORK OPERATIONAL DOMAIN ASSESSMENT (NODA)

Methodology Specification — Version 1.0

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.

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

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

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

Both H and B are computed on a 0–1 scale and then expressed as integers from 0 to 100. 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 within the same Category **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.

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 **MUST NOT** reduce the weight of the Check's Category. The remaining Unconditional and applicable Conditional Checks within that Category continue to carry the Category's full published weight w_i , unredistributed to other Categories.

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.

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.

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%
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); 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, pending additional Checks scheduled for introduction prior to public launch (Section 13). CAT-SEC, formerly single-Check, was expanded to five Checks in a subsequent revision (Section 13) and no longer requires the same near-term attention, though further CAT-SEC Checks remain planned.

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	Likely active service impact or exposure; remediation SHOULD be treated as urgent.
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.

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 inform a Check's existence but do not, by themselves, justify a Critical or Warning Finding, and are instead reflected only in remediation guidance (Section 9). 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), or the reassignment of an already-published Check's Axis or Category. 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 its Axis and Category assigned at introduction, or a Category weight adjustment. MINOR revisions MUST NOT alter the scoring of any Check that existed under the prior version.

— PATCH — editorial clarification with zero scoring impact.

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.

Since this section was first drafted, six additional Checks were introduced under the MINOR-revision discipline of Section 12: 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. 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; CAT-SEC, formerly single-Check, no longer carries the same near-term priority following the additions above.

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

- [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.

[RFC8624] Wouters, P. and Sury, O., "Algorithm Implementation Requirements and Usage Guidance for DNSSEC," June 2019.

[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.

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

Check	Axis (H/B)	Applicability
Parent NS Records	100/0	Unconditional
Nameservers Authoritative	100/0	Unconditional
DNS Parent Sent Glue	100/0	Unconditional
Same Glue	100/0	Unconditional
NS Consistency	100/0	Unconditional
Nameserver Count	100/0	Unconditional
NS TCP	100/0	Unconditional
NS CNAME Check	100/0	Unconditional
NS IPs Public	100/0	Unconditional
SOA Record	100/0	Unconditional
SOA MNAME	100/0	Unconditional
SOA Serial Format	0/100	Unconditional
SOA Serial Sync	100/0	Unconditional
A Record	100/0	Unconditional
A Record IPs Public	100/0	Unconditional
WWW Record	0/100	Unconditional

A.2 CAT-SEC — Security (25%)

Check	Axis (H/B)	Applicability
DNSSEC	0/100	Unconditional
Open Recursor	100/0	Unconditional
AXFR / Zone Transfer	100/0	Unconditional
DNSSEC Multi-Algorithm Signing	0/100	Conditional (DNSSEC)
DNSSEC Algorithm Ban List	0/100	Conditional (DNSSEC)

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

Check	Axis (H/B)	Applicability
SPF Record	100/0	Unconditional — scored even with no MX
SPF Lookup Limit	100/0	Unconditional
DKIM Record	100/0	Conditional (MX)

Check	Axis (H/B)	Applicability
DMARC Record	100/0	Unconditional — scored even with no MX

A.4 CAT-RELI — Reliability (15%)

Check	Axis (H/B)	Applicability
NS Different Subnets	30/70	Unconditional
ASN Diversity	30/70	Unconditional
MX Records	100/0	Unconditional — absence of MX is a valid Pass
MX Validity	100/0	Conditional (MX)
MX Not IP	100/0	Conditional (MX)
MX CNAME	100/0	Conditional (MX)
MX Consistency	100/0	Conditional (MX)
MX Redundancy	30/70	Conditional (MX)
Duplicate MX IPs	100/0	Conditional (MX)
MX IPs Public	100/0	Conditional (MX)
MX Reverse DNS (PTR)	0/100	Conditional (MX)

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

Check	Axis (H/B)	Applicability
Blacklist Check	100/0	Unconditional

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
SOA Refresh	—/—	Unconditional	Infra-adjacent
SOA Retry	—/—	Unconditional	Infra-adjacent
SOA Expire	—/—	Unconditional	Infra-adjacent
SOA Minimum TTL	—/—	Unconditional	Infra-adjacent

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:

$H = 1.2725 / 1.3400 = 0.9496 \rightarrow$ Health Score: 95

$B = 0.5025 / 1.0600 = 0.4740 \rightarrow$ Best Practice Score: 47

Per the interpretation tiers in Section 6, 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 Needs Improvement — 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.