

CMP-001

NetBox and Nautobot

Network source-of-truth platforms



CURRENT EVIDENCE SNAPSHOT

Official product sources refreshed through 2026-09-17

Controlled Mythos laboratory rerun: NOT ASSERTED

VERSION 1.2.0-candidate

CANDIDATE COMPARATIVE EVALUATION / PUBLIC

BENCHMARK DOCTRINE

Gates before scores. Evidence before certainty. Profile fit before universal ranking.

NetBox and Nautobot

Network source-of-truth platforms

DOCUMENT ID

CMP-001

VERSION

1.2.0-candidate

STATUS

Candidate Comparative Evaluation

PUBLICATION DATE

2026-09-17

SCHEDULED REVIEW

2026-12-16

CANONICAL ROUTE

/library/evaluations/cmp-001/

2

CANDIDATES

18

ATOMIC CRITERIA

9

MANDATORY GATES

3

WORKLOAD PROFILES

E2

CURRENT MAX EVIDENCE

CURRENT DECISION BOUNDARY

Do not issue a current winner. The July baseline is a near-tie and both platforms changed materially. Freeze the candidate versions, then rerun upgrade, reconciliation, integration, failure, and export tests before selecting an authority platform.

NO CURRENT UNIVERSAL WINNER IS PUBLISHED FROM THIS REFRESH.

July 24 baseline scores are preserved as lineage and sensitivity evidence, not silently reissued as September laboratory results.

NetBox

v4.7.1 / 2026-09-15

REFRESH TRIGGER

4.7 line adds new infrastructure models and REST/background-processing changes.

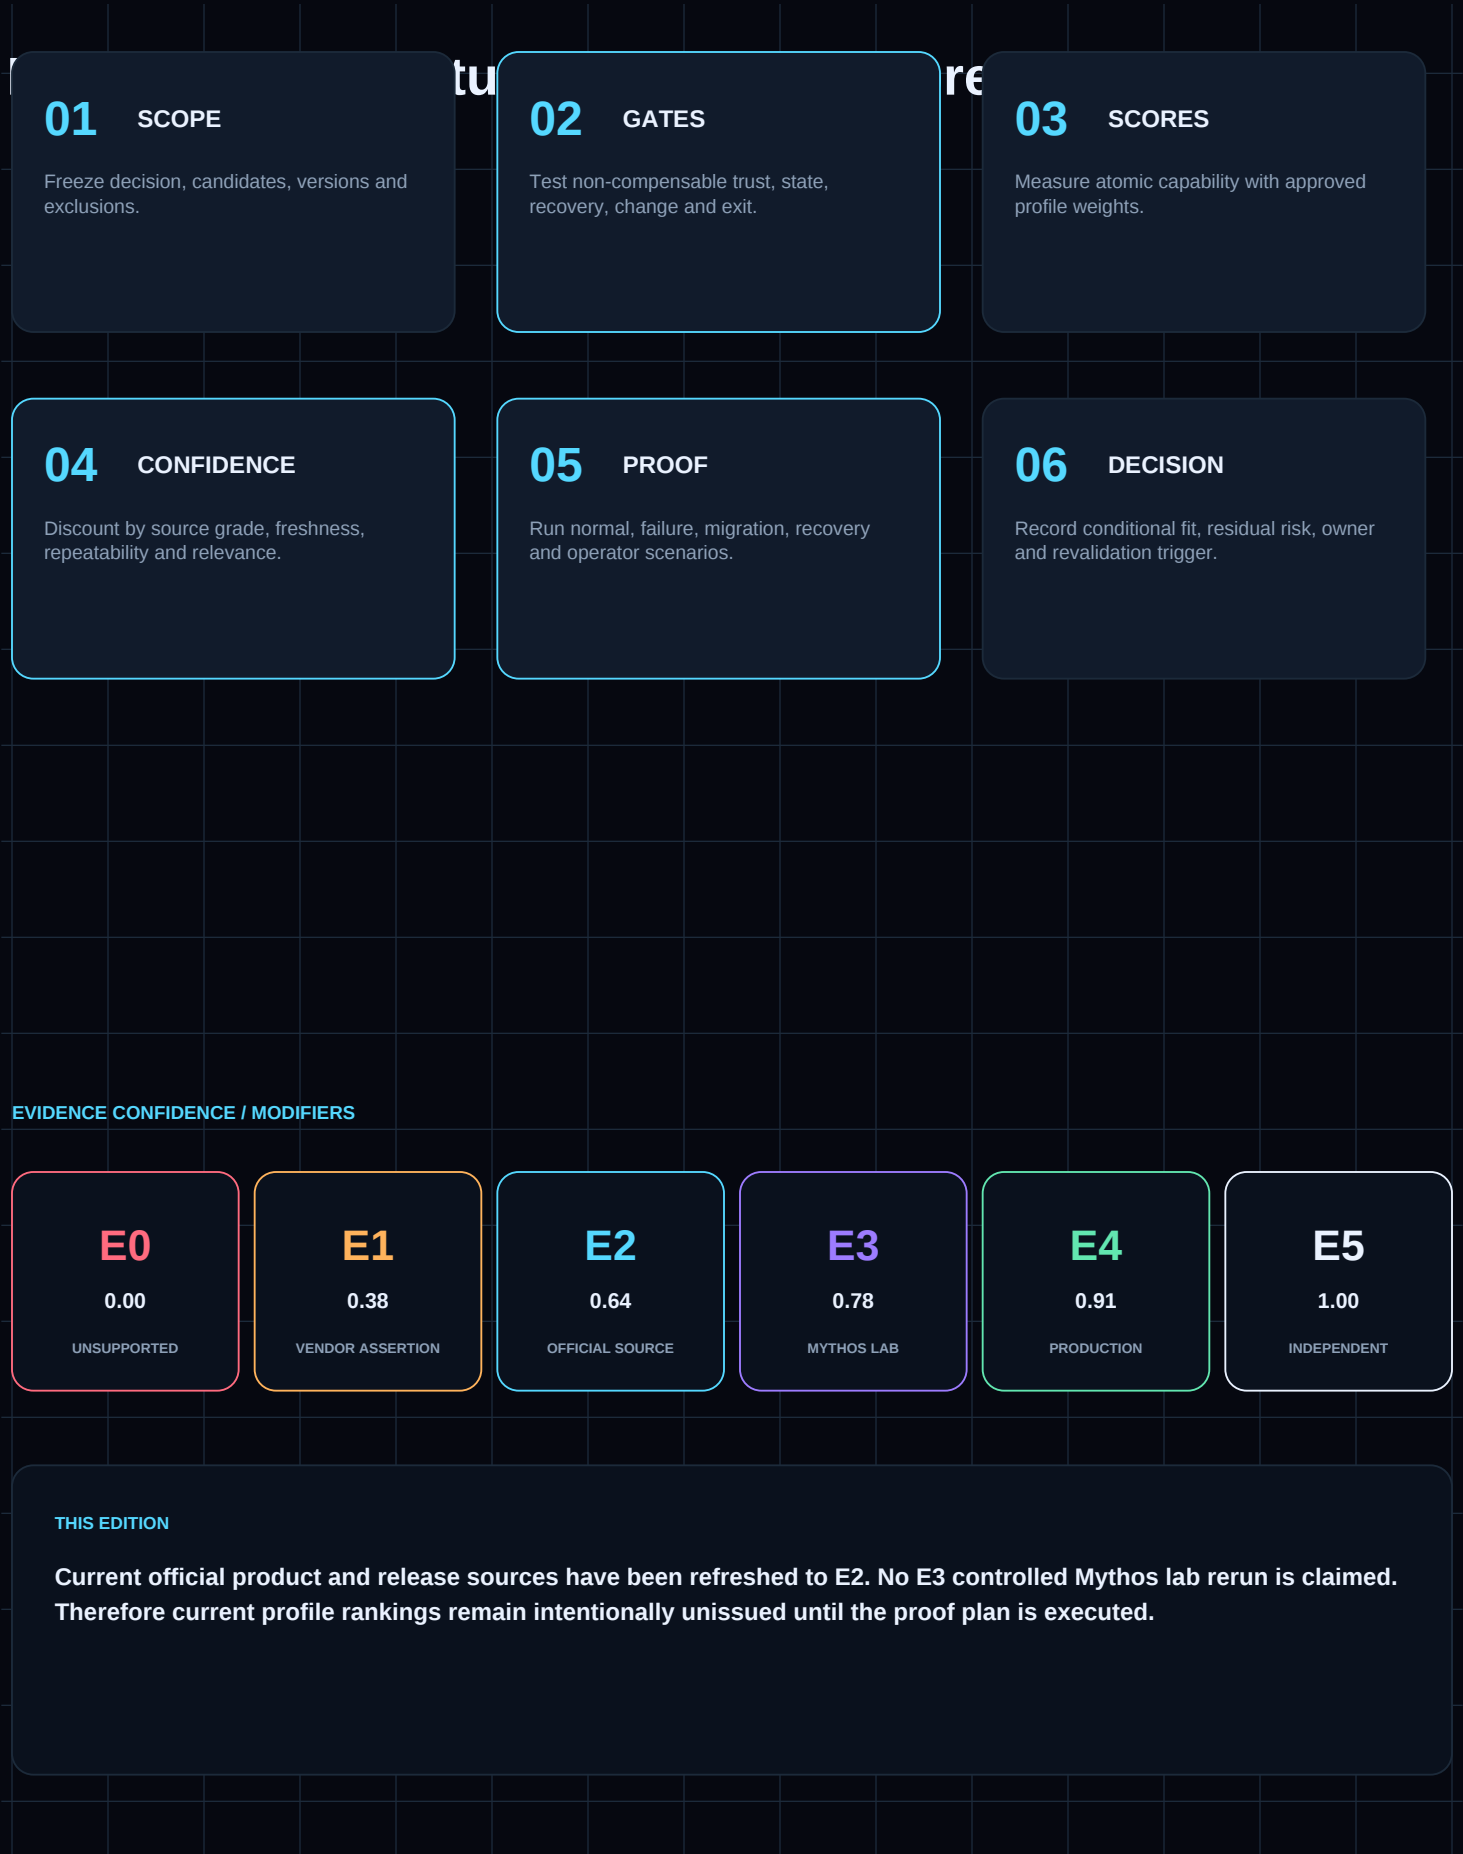
Nautobot

v3.2.5 / 2026-09-14

REFRESH TRIGGER

3.2 includes security fixes and job-execution API behavior changes requiring migration attention.

FINAL BENCHMARK SYSTEM / DECISION ARCHITECTURE



What changed since the July benchmark snapshot

Changes below are sourced from current official release documentation. They identify revalidation triggers; they are not translated into new numeric scores without controlled proof.

NetBox

v4.7.1 / 2026-09-15

MATERIAL DELTA

4.7 line adds new infrastructure models and REST/background-processing changes.

CRITERIA TOUCHED

C01 / C05 / C08 / C14

MODERATE REVALIDATION

Nautobot

v3.2.5 / 2026-09-14

MATERIAL DELTA

3.2 includes security fixes and job-execution API behavior changes requiring migration attention.

CRITERIA TOUCHED

C02 / C05 / C12 / C14

HIGH REVALIDATION

MANDATORY GATES / CURRENT REVALIDATION POSTURE

Mandatory gates remain non-compensable

No current release is marked PASS solely from documentation. E2 means the official source supports the capability path; LAB and RETEST identify proof still required. HOLD blocks a current decision.

MANDATORY GATE	NETBOX	NAUTOBOT
C01 Functional coverage	RETEST	E2
C02 Security / trust	E2	RETEST
C03 Governance / approval	E2	E2
C04 State integrity	E2	E2
C07 Observability / evidence	E2	E2
C09 Resilience / continuity	LAB	LAB
C11 Change safety / rollback	LAB	LAB
C16 Portability / exit	LAB	LAB
C18 Assurance confidence	HOLD	HOLD

LEGEND

E2

official-source support only

LAB

controlled proof required

RETEST

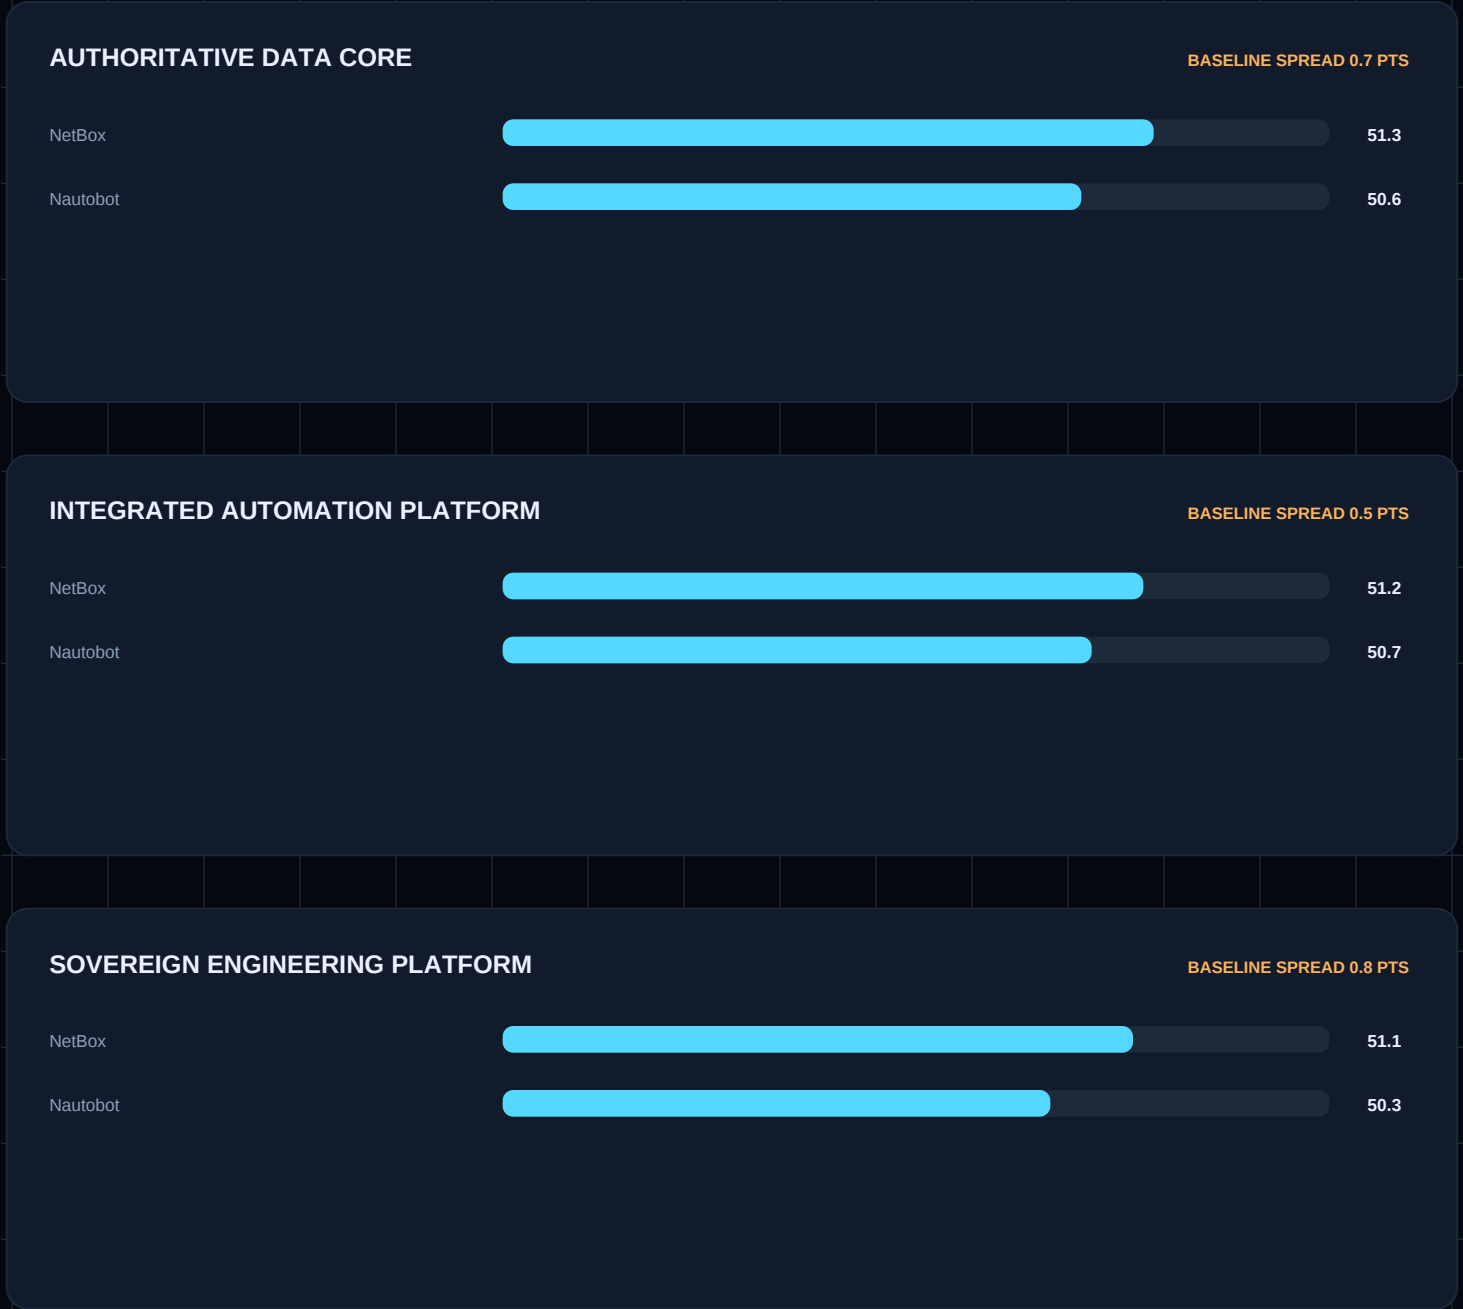
release delta touches this gate

HOLD

decision gate remains closed

Historical profile sensitivity - preserved, not reissued

Values below are the 2026-07-24 evidence-adjusted baseline. They show sensitivity to operating-model weights. The September refresh does not recompute rank because E3 lab proof has not been reproduced.



CURRENT RANK STATUS: WITHHELD PENDING CONTROLLED PROOF.

Evidence confidence is separate from capability

E5

Independent repeatable validation

confidence modifier 1.00

E4

Production evidence in adopting environment

confidence modifier 0.91

E3

Controlled Mythos laboratory result

confidence modifier 0.78

E2

Official documentation / release / source

confidence modifier 0.64

E1

Vendor assertion or marketing

confidence modifier 0.38

E0

Unsupported or unverified

confidence modifier 0.00

CURRENT EVIDENCE STATE

NetBox

E2 REFRESHED

Historical adjusted confidence: 59.7%

E3 LAB: NOT REPRODUCED

Nautobot

E2 REFRESHED

Historical adjusted confidence: 59.7%

E3 LAB: NOT REPRODUCED

ASSURANCE RULE

A current release note can change capability context, but it cannot raise assurance beyond E2. Current recommendation strength is therefore bounded until the test plan produces reproducible artifacts.

MIGRATION, LIFECYCLE AND IRREVERSIBLE-COMMITMENT RISK

Lifecycle risk is evaluated independently of features

The benchmark models migration, upgrade, compatibility, support, staffing, data/state export, downtime and exit. Current release changes below are explicit proof triggers.

NetBox

MODERATE REVALIDATION

4.7 line adds new infrastructure models and REST/background-processing changes.

RETEST: C01 / C05 / C08 / C14

Nautobot

HIGH REVALIDATION

3.2 includes security fixes and job-execution API behavior changes requiring migration attention.

RETEST: C02 / C05 / C12 / C14

IRREVERSIBLE COMMITMENT GATE

Before migration or procurement lock-in, prove rollback, state portability, operational reconstruction, and supplier-exit assumptions.

Controlled proof is the next decision-producing step

REPRODUCTION STATE / NOT EXECUTED IN THIS EVIDENCE-REFRESH RELEASE

01

CONTROLLED SCENARIO

Model a representative multi-site network with IPAM, circuits, devices, virtualization, tenancy, and custom data.

02

CONTROLLED SCENARIO

Import authoritative and conflicting data from two systems and prove deterministic reconciliation.

03

CONTROLLED SCENARIO

Execute a guarded change workflow with approval, dry-run, evidence, and rollback.

04

CONTROLLED SCENARIO

Measure API throughput, query latency, job concurrency, and database recovery.

05

CONTROLLED SCENARIO

Upgrade a realistic plugin or application set and record compatibility failures.

06

CONTROLLED SCENARIO

Export the full data model and rebuild the environment from controlled backups.

EVIDENCE PACKAGE

Preserve exact versions, architecture, configuration, scenario, expected result, instrumentation, raw output, timing, failures, operator actions, deviations, acceptance state, owner, and artifact hashes.

Current decision posture

Do not issue a current winner. The July baseline is a near-tie and both platforms changed materially. Freeze the candidate versions, then rerun upgrade, reconciliation, integration, failure, and export tests before selecting an authority platform.

CURRENT RANKING: NOT REISSUED / CONTROLLED LAB REQUIRED

CURRENT OFFICIAL-SOURCE REGISTER / CHECKED 2026-09-17

NetBox

E2 / OFFICIAL

NetBox GitHub releases

github.com

Nautobot

E2 / OFFICIAL

Nautobot v3.2 release notes

docs.nautobot.com

REVISION LINEAGE

- 1.2.0-candidate / 2026-09-17 - current-source evidence refresh; no lab rescore issued.
- 1.1.0-candidate / 2026-07-24 - baseline benchmark using the final comparative evaluation system.
- Trigger earlier review after major release, acquisition, licensing change, critical vulnerability, material outage, failed PoC, or workload change.

BASELINE ANALYTIC RECORD CONTINUES ON PAGE 11 FOR TRACEABILITY.

ATOMIC CRITERIA MODEL

WEIGHTED CAPABILITY WITH EXPLICIT MINIMUM EVIDENCE

ID	CRITERION	WEIGHT	GATE	MINIMUM EVIDENCE
C01	Network source-of-truth functional coverage	5	YES	Feature documentation, APIs, configuration exports, and scenario demonstration.
C02	Security and trust architecture	5	YES	Threat model, identity flows, RBAC tests, audit records, and security configuration.
C03	Governance and approval controls	5	YES	Approval workflow, policy controls, separation-of-duties tests, and decision records.
C04	Data-model integrity and reconciliation	5	YES	Backup, restore, rollback, drift, and corruption-recovery evidence.
C05	Automation and API quality	4	NO	API specifications, authentication tests, rate limits, event samples, and automation runs.
C06	Integration ecosystem	4	NO	Supported integrations, connector tests, failure behavior, and lifecycle ownership.
C07	Observability and evidence	4	YES	Logs, traces, metrics, reports, export tests, and evidence-retention controls.
C08	Object, API, and job scale	4	NO	Controlled load tests, topology scale, saturation behavior, and capacity model.
C09	Resilience and continuity	5	YES	Failure injection, HA tests, recovery timing, degraded-mode operation, and runbooks.
C10	Usability and operator cognition	3	NO	Task observation, error-rate measures, navigation tests, and operator interviews.
C11	Change safety and rollback	5	YES	Plan or preview output, canary tests, rollback execution, and post-change verification.
C12	Extensibility and programmability	3	NO	Plugin, module, provider, workflow, or code-extension tests and upgrade impact analysis.
C13	Deployment and control-plane sovereignty	3	NO	Deployment architecture, data-flow mapping, residency evidence, and offline tests.
C14	Lifecycle and upgrade discipline	4	NO	Release notes, upgrade test, compatibility matrix, support policy, and migration plan.
C15	Economics and cost predictability	3	NO	Bill-of-materials, licensing model, staffing estimates, metering, and sensitivity analysis.
C16	Portability and exit	4	YES	Export tests, format analysis, replacement workflow, and decommission evidence.
C17	Vendor and ecosystem risk	3	NO	License analysis, roadmap evidence, bus-factor indicators, and dependency inventory.
C18	Assurance confidence	5	YES	Source provenance, lab artifacts, production evidence, independent replication, and review date.

SCORE SCALE

0 absent; 1 materially deficient; 2 partial; 3 adequate with limits; 4 strong; 5 leading for the defined profile. Scores remain provisional until the required evidence grade is achieved.

RAW CAPABILITY SCORECARD

DOCUMENTED CAPABILITY BEFORE EVIDENCE DISCOUNT

CRITERION	NETBOX	NAUTOBOT
C01 Network source-of-truth functional coverage	4.6	4.5
C02 Security and trust architecture	4.1	4.0
C03 Governance and approval controls	3.8	4.0
C04 Data-model integrity and reconciliation	4.4	4.2
C05 Automation and API quality	4.7	4.6
C06 Integration ecosystem	4.4	4.5
C07 Observability and evidence	4.1	4.2
C08 Object, API, and job scale	4.3	4.0
C09 Resilience and continuity	4.0	3.9
C10 Usability and operator cognition	4.2	4.1
C11 Change safety and rollback	4.0	4.2
C12 Extensibility and programmability	4.4	4.8
C13 Deployment and control-plane sovereignty	4.7	4.5
C14 Lifecycle and upgrade discipline	4.2	4.0
C15 Economics and cost predictability	4.5	4.1
C16 Portability and exit	4.6	4.3
C17 Vendor and ecosystem risk	4.4	4.1
C18 Assurance confidence	3.5	3.5

WEIGHTED BASELINE / 100

NetBox

84.8

Nautobot

83.4

EVIDENCE-ADJUSTED SCORECARD

CAPABILITY DISCOUNTED BY CURRENT EVIDENCE CONFIDENCE

CANDIDATE	RAW	EVIDENCE CONFIDENCE	ADJUSTED	CURRENT STATE
NetBox	84.8	59.7%	50.7	CONTROLLED LAB PENDING
Nautobot	83.4	59.7%	50.0	CONTROLLED LAB PENDING

EVIDENCE SCALE

E0 / 0.00

Unsupported or unverified

E1 / 0.38

Vendor assertion or marketing statement

E2 / 0.64

Official documentation, release notes, or source repository

E3 / 0.78

Controlled Mythos laboratory result

E4 / 0.91

Production evidence from the adopting environment

E5 / 1.00

Independent repeatable validation

CURRENT CONFIDENCE BOUNDARY

Most candidate criteria are supported at E2: official documentation or source evidence. Environment-specific laboratory, production, and independent evidence remain future gates.

PROFILE-SPECIFIC RANKINGS

EVIDENCE-ADJUSTED SENSITIVITY ANALYSIS

AUTHORITATIVE DATA CORE

Prioritizes clean network modeling, external-tool interoperability, portability, and low coupling.

1. NetBox	51.3
2. Nautobot	50.6

INTEGRATED AUTOMATION PLATFORM

Prioritizes embedded execution, extensibility, synchronization, workflow governance, and operator productivity.

1. NetBox	51.2
2. Nautobot	50.7

SOVEREIGN ENGINEERING PLATFORM

Prioritizes self-hosting, open interfaces, exportability, and control over data and runtime dependencies.

1. NetBox	51.1
2. Nautobot	50.3

RANK SENSITIVITY

Small score differences are not decision certainty. Treat near-ties as a requirement for deeper PoC, commercial, migration, and operating-model evidence.

CANDIDATE DEEP DIVE / 01

NETBOX

OPERATING MODEL

Focused network source-of-truth platform centered on DCIM, IPAM, extensible data modeling, APIs, plugins, and integration with external automation systems.


DOCUMENTED STRENGTHS

- Mature network modeling and documentation posture.
- Strong REST API and broad ecosystem adoption.
- Clear separation between source-of-truth data and external execution systems.
- Open-source core with commercial support paths.

CAUTIONS AND PROOF OBLIGATIONS

- Complex orchestration generally remains external.
- Plugin governance and upgrade compatibility require discipline.
- Desired-state workflows must be engineered around the authoritative model.

BASELINE SCORE PROFILE



GATE DISPOSITION

PASS - core source-of-truth profile | CONDITIONAL - embedded workflow profile requires external orchestration

CANDIDATE DEEP DIVE / 02

NAUTOBOT

OPERATING MODEL

Network automation platform combining source-of-truth data, jobs, applications, design workflows, and synchronization patterns.


DOCUMENTED STRENGTHS

- Integrated Jobs and application framework.
- Strong desired-state and automation-platform orientation.
- SSoT synchronization patterns and DiffSync ecosystem.
- Extensible workflows can reduce handoffs between data and execution.

CAUTIONS AND PROOF OBLIGATIONS

- Integrated execution increases platform responsibility and governance scope.
- Application compatibility and lifecycle must be controlled.
- Operating model can be heavier than a narrowly scoped source of truth.

BASELINE SCORE PROFILE



GATE DISPOSITION

PASS - integrated automation profile | CONDITIONAL - operating complexity and app lifecycle must be accepted

CROSS-CANDIDATE CONTRASTS

WHERE ARCHITECTURAL DIFFERENCES MATTER MOST

CONTROL-LOOP MODEL

How authority moves from intent to plan, approval, execution, observation, reconciliation, and recovery.

STATE OWNERSHIP

Which system owns desired state, operational state, evidence, history, and conflict resolution.

MANAGED VERSUS SOVEREIGN BOUNDARY

Which runtime, identity, data, policy, relay, or service dependencies remain outside organizational control.

EXTENSION MODEL

Plugins, providers, modules, applications, workflows, SDKs, APIs, and custom code introduce different lifecycle risks.

FAILURE TRANSPARENCY

Candidates differ in how clearly they expose partial execution, degraded mode, retries, inconsistency, and recovery state.

EXIT MECHANICS

Export formats, state semantics, policy portability, knowledge transfer, and replacement effort are materially different.

CANDIDATE	CURRENT ADVANCEMENT POSITION
NetBox	PASS - core source-of-truth profile
Nautobot	PASS - integrated automation profile

CROSS-CANDIDATE CONTRASTS

WHERE ARCHITECTURAL DIFFERENCES MATTER MOST

CONTROL-LOOP MODEL

How authority moves from intent to plan, approval, execution, observation, reconciliation, and recovery.

STATE OWNERSHIP

Which system owns desired state, operational state, evidence, history, and conflict resolution.

MANAGED VERSUS SOVEREIGN BOUNDARY

Which runtime, identity, data, policy, relay, or service dependencies remain outside organizational control.

EXTENSION MODEL

Plugins, providers, modules, applications, workflows, SDKs, APIs, and custom code introduce different lifecycle risks.

FAILURE TRANSPARENCY

Candidates differ in how clearly they expose partial execution, degraded mode, retries, inconsistency, and recovery state.

EXIT MECHANICS

Export formats, state semantics, policy portability, knowledge transfer, and replacement effort are materially different.

CANDIDATE	CURRENT ADVANCEMENT POSITION
NetBox	PASS - core source-of-truth profile
Nautobot	PASS - integrated automation profile

ARCHITECTURE AND INTEGRATION FIT

THE PRODUCT IS ONLY ONE PART OF THE OPERATING SYSTEM

REFERENCE OPERATING LAYERS

01 Authoritative data model and ownership boundaries

02 API and event integration layer

03 Validation, reconciliation, and drift workflow

04 Execution platform or external orchestrators

05 Evidence, audit, and recovery store

INTEGRATION BOUNDARIES

- Identity provider and RBAC
- Git and CI/CD
- Telemetry and event bus
- ITSM and change management
- Device discovery and collectors
- Ansible, Terraform/OpenTofu, and network automation frameworks

ARCHITECTURE GATE

Every external dependency requires an owner, identity boundary, failure mode, evidence path, version policy, recovery procedure, and exit strategy.

SECURITY, OPERATIONS, LIFECYCLE, ECONOMICS, AND EXIT

CROSS-DOMAIN DECISION RECORD

SECURITY AND AUTHORITY

Identity, credentials, secrets, roles, privileged actions, signing, approvals, isolation, audit, and incident response must be tested as an integrated system.

RELIABILITY AND RECOVERY

Control-plane, data-plane, dependency, region, database, queue, network, and operator failures require defined degraded modes and recovery objectives.

CHANGE AND LIFECYCLE

Version, compatibility, migration, canary, rollback, content, plugin, provider, firmware, and decommission procedures are part of product fitness.

ECONOMICS AND CAPACITY

Licenses, metering, data, compute, hardware, storage, support, staffing, training, integration, migration, and opportunity cost require sensitivity analysis.

SOVEREIGNTY AND PRIVACY

Data location, support access, telemetry, subprocessors, encryption, key control, disconnected operation, and legal obligations must align with policy.

PORTABILITY AND EXIT

Configuration, policy, data, state, evidence, workflows, identities, and operational knowledge must be exportable and reconstructable.

DECISION OWNERSHIP

Architecture, security, operations, finance, procurement, legal, data, and service owners jointly accept the final profile, evidence, residual risk, and exit obligations.

RISK AND FAILURE REGISTER

SCENARIOS THAT CAN INVALIDATE A FEATURE-BASED SELECTION

ID	SEVERITY	FAILURE SCENARIO	REQUIRED TREATMENT
R-01	CRITICAL	A source of truth becomes a passive inventory because no authority model is enforced.	PREVENT / DETECT / RECOVER / EVIDENCE
R-02	HIGH	Multiple systems write overlapping fields without reconciliation ownership.	PREVENT / DETECT / RECOVER / EVIDENCE
R-03	HIGH	Plugins or applications create upgrade deadlocks.	PREVENT / DETECT / RECOVER / EVIDENCE
R-04	HIGH	Automation consumes stale or incomplete data and produces unsafe changes.	PREVENT / DETECT / RECOVER / EVIDENCE
R-05	MEDIUM	Backup restores records but not integrations, jobs, secrets, or evidence.	PREVENT / DETECT / RECOVER / EVIDENCE
R-06	MEDIUM	A platform is selected for feature breadth without a sustainable operating model.	PREVENT / DETECT / RECOVER / EVIDENCE

RISK RULE

A candidate with an unresolved critical failure path cannot advance because optional capability, market position, or commercial discount cannot compensate for missing safety or recovery evidence.

CONTROLLED PROOF-OF-CAPABILITY PLAN

REPEATABLE SCENARIOS, INSTRUMENTATION, AND ACCEPTANCE

TEST 01

Model a representative multi-site network with IPAM, circuits, devices, virtualization, tenancy, and custom data.

TEST 02

Import authoritative and conflicting data from two systems and prove deterministic reconciliation.

TEST 03

Execute a guarded change workflow with approval, dry-run, evidence, and rollback.

TEST 04

Measure API throughput, query latency, job concurrency, and database recovery.

TEST 05

Upgrade a realistic plugin or application set and record compatibility failures.

TEST 06

Export the full data model and rebuild the environment from controlled backups.

EVIDENCE PACKAGE

Preserve versions, architecture, configuration, data set, scenario, expected result, instrumentation, raw output, timing, failures, operator actions, deviations, acceptance state, owner, and artifact hashes.

CONDITIONAL RECOMMENDATION AND REVALIDATION

DECISION RECORD, ACCEPTANCE, AND NEXT EVIDENCE

RECOMMENDATION POSITION

- Select NetBox for a focused data authority with external deterministic execution.
- Select Nautobot for an integrated network automation platform with governed application lifecycle.
- Require an environment-specific PoC before committing the authoritative data model.
- Revalidate releases, plugins, commercial packaging, and operational scale every 90 days.

CANDIDATE	ADJ. SCORE	GATE POSITION	LAB STATE
NetBox	50.7	PASS - core source-of-truth profile; CONDITIONAL - embedded workflow profile requires external orchestration	PENDING
Nautobot	50.0	PASS - integrated automation profile; CONDITIONAL - operating complexity and app lifecycle must be accepted	PENDING

REVALIDATION SCHEDULE

Review no later than 2026-10-24. Review earlier after a major release, commercial change, critical vulnerability, material outage, architecture transition, failed acceptance test, or change to the target workload profile.

SOURCE AUTHORITY AND REVISION

CURRENT EVIDENCE SNAPSHOT AND PUBLICATION HISTORY

SOURCE ID	PUBLISHER	TITLE	CHECKED
NIST-CSF-20	NIST	Cybersecurity Framework 2.0	2026-07-24
NIST-SP800-53	NIST	Security and Privacy Controls for Information Systems and Organizations	2026-07-24
NETBOX-DOCS	NetBox Labs	NetBox Documentation	2026-07-24
NETBOX-API	NetBox Labs	NetBox REST API Overview	2026-07-24
NETBOX-GH	NetBox Community	NetBox Source Repository	2026-07-24
NAUTOBOT-DESIGN	Network to Code	Nautobot Design Philosophy	2026-07-24
NAUTOBOT-JOBS	Network to Code	Nautobot Jobs Developer Guide	2026-07-24
NAUTOBOT-SSOT	Network to Code	Nautobot Single Source of Truth Application	2026-07-24

SOURCE POSITION

Official documentation and source repositories support the current capability model. Source records preserve publisher, title, URL, purpose, and review date in the machine-readable authority register. Commercial terms, performance, and environment-specific behavior remain unverified until controlled proof.

REVISION HISTORY

1.1.0-candidate / 2026-07-24 - permanent-publication rebuild using the final benchmark system, profile-specific rankings, evidence adjustment, mandatory gates, and controlled PoC requirements.