NETBRAIN:

Agentic AI for Network Automation

NetBrain's next-gen platform delivers intelligent automation that works like an extension of your team—combining advanced AI with rigorous governance to ensure every action is accurate, attributable, and auditable. According to Gartner, "This is the biggest shift in network operations in ~20 years." By 2030, agent-initiated execution will become the primary approach for network runtime activities. NetBrain is purpose-built for this new era, moving beyond static dashboards to a world where humans and agents collaborate seamlessly through deterministic, constrained intelligence.

The Problem AI Is Actually Solving

We address the critical gap in network operations: not a lack of tools, but inconsistent execution under pressure. Across enterprises and service providers, network operations fail in consistent and well-understood ways:

- Troubleshooting remains manual and heavily dependent on individual experience, making expertise a single point of failure.
- Diagnostic approaches vary significantly between engineers, leading to inconsistent outcomes and finger-pointing during outages.
- Root cause analysis quality degrades under pressure, with critical steps missed when they matter most.
- Changes are executed reactively with limited validation, introducing risk rather than reducing it.
- Operational evidence is captured after incidents, not during them, forcing teams to reconstruct what happened rather than knowing for certain.
- Knowledge remains tribal rather than institutional, walking out the door when engineers leave and forcing the next team to start from scratch.

These are not governance failures. Policies, processes, and tooling usually exist. The failure occurs at execution time, when complexity and urgency overwhelm manual workflows.

Hybrid, cloud, and third-party environments make this worse. Visibility becomes inconsistent. Tooling is often vendor specific. Cross-domain diagnosis depends on human correlation across platforms, technologies, and teams—increasing both the time required to reach a conclusion and the chance of getting it wrong.

NetBrain applies AI to address this execution problem directly standardizing how work gets done, not just what work is tracked. Its value is in standardizing good operational practice and making that discipline available consistently across the team.

Why AI Must Be Bounded

As AI adoption accelerates, a new risk profile emerges alongside the problems it solves. The risk is not AI in itself—it is unbounded AI.

When AI drives actions without clear authority boundaries, it can introduce unintended changes, widen blast radius, and accelerate failures before humans can intervene. When reasoning is opaque, decisions become impossible to audit or defend. When actions are not logged, evidence cannot be reconstructed.

That is not resilience. It is a new failure mode.

NetBrain AI is designed to strengthen operational discipline, not weaken it—by keeping authority bounded, evidence preserved, and human accountability intact.

How NetBrain Approaches AI

We apply domain-specific, operationally constrained AI to standardize execution, accelerate troubleshooting, and produce trusted operational evidence, while keeping authority and accountability firmly under human control. Unlike probabilistic black boxes, NetBrain delivers deterministic outcomes—every action is traceable, repeatable, and grounded in verified data.

PRINCIPLE 1

Operational Truth Over Inference

AI conclusions are grounded in executed diagnostics, digital twin data, and time-stamped results—never invented explanations or probabilistic assumptions.

PRINCIPLE 2

Intelligence Without Unbounded Authority

AI may investigate, diagnose, and recommend. It may not execute remediation, bypass change control, or operate outside defined permissions.

PRINCIPLE 3

Human Accountability Preserved

Every AI-assisted action is attributable to a user, bound to a workflow, and governed by role-based permissions. AI accelerates—it does not replace.

How NetBrain Complements Your Existing Tools

Vendor-native AI is strongest where telemetry is homogeneous, and tooling is tightly coupled to the platform.

NetBrain operates differently. It complements vendor-native capabilities by working across the execution, reasoning, and governance layer in multi-vendor environments.

It is not trying to replace vendor analytics, hardware-specific insight, or security tooling. It provides cross-vendor operational reasoning, validation, workflow-driven diagnosis, and defensible evidence across environments that are typically fragmented by technology, team, and operational boundary.

Most real-world networks are not operationally neat. They are mixed, layered, and dependent on consistent cross-domain execution—which is where NetBrain is designed to operate.

The Three Pillars of NetBrain AI

By unifying three transformative AI technologies within a deterministic, governed framework, NetBrain reduces execution variability and mean-time-to-resolution (MTTR) without increasing operational risk.

- Constrained Agentic AI serves as a digital investigator, autonomously executing complex, multi-vendor troubleshooting tasks. It decides what to investigate and executes diagnostics, but it is explicitly separated from authority—it does not perform autonomous remediation. Every decision follows deterministic guardrails.
- Generative AI acts as your always-available analyst, translating technical CLI outputs and automation results into plain-English insights for faster handovers and clearer communication.
- Retrieval-Augmented Generation (RAG) ensures every insight and recommendation is grounded in your actual network documentation, configurations, and runbooks—not generic internet best practices. This grounding creates deterministic, auditable results.

This powerful combination delivers the industry's first context-aware and governed network automation solution that:

- Understands your specific environment by accessing real-time device documentation and configuration databases.
- Learns from both live network data and your institutional knowledge.
- Acts with precision by grounding every decision in verified data, eliminating the "black box" anxiety of AI hallucinations.
- Preserves accountability by keeping all AI actions attributable to a user, bound by workflows, and governed by role-based permissions.
- Delivers deterministic outcomes through traceable reasoning chains and repeatable execution paths.

The NetBrain AI Advantage: Truth Over Inference

Unlike conventional solutions that rely on probabilistic inference, NetBrain's architecture ensures operational truth:

- Decisions are grounded in executed diagnostics, not assumed patterns.
- Responses are tied to your network's actual documentation, reflecting current best practices.
- Intelligence is separated from authority; AI may investigate and recommend, but it cannot bypass change control.
- Full traceability provides a time-stamped reasoning chain for every investigation, satisfying audit and regulatory scrutiny.
- Repeatable outcomes ensure the same input consistently produces the same verified result.

The result is enterprise-grade automation you can trust, reducing mean-time-to-resolution while strengthening execution discipline.

Validated Against the Gartner BRITE Framework

True agentic behavior requires more than just responding to prompts. It requires event-driven initiation, explainable planning, and governed execution. NetBrain's architecture is built to meet the rigorous standards of the Gartner BRITE framework:

BRITE Element	Gartner Definition	NetBrain Delivery
Baseline	Show intended and observed state; use deltas for planning and verification.	Real-time digital twin maintains continuous observed state across multi-vendor environments, with intended state derived from golden configs and operational learning.
Reason	Convert goals or signals into explainable multi-step plans with rationale traces.	Deep Diagnosis dynamically builds investigation paths, interprets results, and produces explicit, time-stamped reasoning chains for every action.
Integrate	Invoke existing tools and APIs; export artifacts to ITSM and related systems.	Native integrations with ServiceNow, Splunk, and collaboration tools; all actions and artifacts are logged and exportable.
Trigger	Initiate from events, policies, or thresholds—not just prompts.	Triggered AI automation auto-runs diagnostics from incidents, anomalies, and policy violations without human initiation.
Effect	Produce verifiable state changes with execution logs and rollback.	Every diagnostic produces verifiable evidence; change workflows include pre/post-checks and one-click rollback.

AI-Driven Observability & Intelligent Troubleshooting

NetBrain transforms network operations through an AI-powered closed-loop system that unifies:

1. Comprehensive Observability

- Real-time topology mapping with dependency analysis.
- Continuous baseline establishment for all network elements.
- Flow-based traffic visualization and anomaly detection.

2. Intelligent Assessment

- Automated health scoring across devices, paths, and services.
- Proactive risk identification through ML-powered pattern recognition.
- Impact analysis for planned changes using digital twin modeling.

3. Context-Aware & Constrained Troubleshooting

Deep Diagnosis (Agentic Reasoning)

The most advanced AI capability in NetBrain. Deep Diagnosis does not follow a static script. It:

- Accepts an operational problem statement
- Determines which diagnostics are relevant
- Executes diagnostics using NetBrain automation
- Interprets results and key findings
- Adjusts investigation paths dynamically
- Iterates until sufficient evidence exists
- Produces a conclusion with an explicit reasoning chain

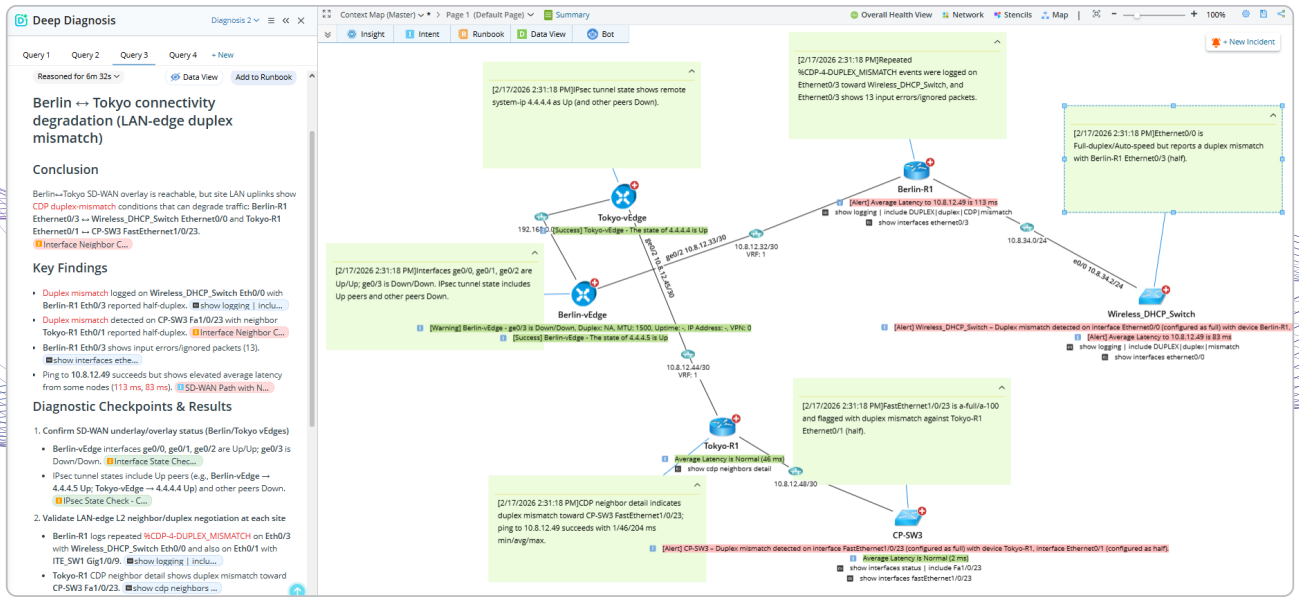

Each step is executed, recorded, and traceable. This qualifies as agentic reasoning because the system makes investigative decisions dynamically and adapts based on outcomes. It is deliberately constrained and does not perform autonomous remediation.

Deep Diagnosis can only operate within:

- Existing NetBrain automations
- Network-scoped data
- Defined permissions and workflows

It cannot invent commands, bypass controls, or act independently of governance.

Operational impact includes:

- Standardized troubleshooting quality
- Reduced dependency on senior engineers
- Improved MTTR without increased risk
- Defensible investigation evidence
- Conversion of investigations into reusable runbooks

Guided Diagnostics
Uses institutional runbooks and RAG architecture to recommend next steps without taking uncontrolled action.

Automated Remediation Playbooks
Generated with validation checks, integrated with ITSM workflows for approval before execution.

Validation, Assurance, and Change Safety

Beyond Deep Diagnosis, NetBrain includes additional AI-powered validation tools that strengthen operational discipline:

- **Quick Assessment:** Runs targeted validation checks against known operational conditions to identify drift, instability, or misalignment early—before they become incidents.
- **TAF (Triggered Automation Framework):** Automates diagnostics and validation based on defined triggers and conditions, ensuring consistent execution across the team.
- **Pre/Post-Change Validation:** Changes can be validated before AND after execution, with results interpreted quickly and consistently. This reduces change-related risk and provides immediate confirmation of success or failure.

Issues can be investigated and validated in minutes rather than hours because diagnostics, validation, and evidence gathering are executed as part of the workflow itself.

Key AI Capabilities & Governance

NetBrain's AI capabilities span multiple paradigms, each with distinct roles and constraints.

Agentic AI (The Digital Investigator)

- Executes CLI commands across multi-vendor devices to retrieve real-time status.
- Reads automation results and triggers follow-up diagnostics.
- Creates dynamic dashboards visualizing performance and faults.
- Retrieves device properties: IP lookup, L2/L3 neighbour discovery, and more.
- **Constraint:** Operates only within existing automation libraries and defined permissions. Cannot invent commands or execute changes. Every action is deterministic and auditable.

Interpretive & Analytical AI (The Communication Layer)

This class of AI focuses on understanding and communication rather than execution. It reads structured execution outputs, explains findings, and supports communication with non-network stakeholders. It does not initiate execution or modify workflows.

- **AI Runbook Companion:** Provides step-by-step guidance during incidents by surfacing relevant runbooks and institutional knowledge.
- **AI Document Generation:** Automatically creates and updates network documentation based on live network state.
- **Ticket-Triggered Alerting Framework:** When integrated with ITSM tools, generates plain-English summaries of diagnostic findings attached to tickets.
- **Insight Analysis:** Analyzes automation results and explains issues in plain English, clarifying why alerts were triggered.

The result is faster handovers, clearer incident communication, and improved post-incident documentation—all without increasing operational risk.

Generative AI (The Natural Language Interface)

- Processes raw CLI outputs into human-readable insights.
- Answers natural language queries (e.g., "Show VPN tunnels with latency >100ms").
- Explains automation outcomes and the "why" behind alerts.
- **Constraint:** Explains executed data; does not initiate execution.

Parser Assistant (Human-in-the-Loop)

Network automation depends on structured data, yet network devices produce unstructured output. Manually building parsers is time-consuming, error-prone, and highly specialized.

- Analyzes sample device output and generates draft parsing logic within defined limits.
- All outputs require explicit human review and acceptance before deployment.
- Cuts parser development time from hours to minutes while maintaining human oversight.
- **Constraint:** Cannot deploy parsers automatically, override inheritance logic, or bypass validation controls. Preserves human accountability while accelerating parser development.

AI Governance & Security

Data Privacy & Security

- **Data Privacy:** Customer data is never used to train public LLMs. Formal confirmation available in customer documentation.
- **Encryption:** AES-256 for data at rest and in transit.
- **Access Control:** Role-Based Access Control (RBAC) and Zero Trust principles ensure human accountability.
- **Session Scoping:** AI interactions are session-scoped only; minimum required data transmitted per request.
- **Auditability:** Every AI interaction and configuration change is logged with timestamps and full user attribution.

Compliance

- **Regulatory Alignment:** Supports GDPR, HIPAA, and SOC 2 Type II.
- **EU AI Act (Reg. 2024/1689):** Aligns with limited-risk category by design—not used for biometric ID, no direct individual impact, not a safety-critical autonomous system.
- **Customer-Managed Keys:** Support for customer-provided encryption keys alongside NetBrain-provided keys.
- **Full Disablement:** AI features can be disabled entirely where policy, security review, or regulatory posture requires it.

AI Settings Framework

All AI behavior is configured at domain level—no uncontrolled drift.

Control	Function
Model Role Separation	Assign distinct models by function: general language, reasoning-heavy tasks, complex analysis, and embedding. Each model has a defined purpose.
Embedding Governance	Changes to embedding models trigger automatic reprocessing, preserving determinism and traceability over time.
Controlled Scope	Each AI feature independently enabled or disabled at domain level.
Safe Baselines	Reset AI behavior to known-good configurations at any time.
Role-Based Access	Explicit user privileges control who can initiate AI-assisted operations.
Audit Logging	All AI interactions logged with timestamps and full user attribution.

Knowledge Capture: From Incidents to Institutional Memory

One of the biggest operational problems is not just resolving incidents—it is failing to retain the logic used to resolve them.

NetBrain allows AI-assisted investigations to be converted into governed, versioned runbooks. The output does not disappear when the incident ends. It becomes part of the organization's operating memory.

- Reduces dependence on tribal knowledge
- Improves ramp-up for newer engineers
- Strengthens execution consistency over time
- Makes the operating model more resilient

Practical Applications & Use Cases

1. Automated, Versioned Documentation

Continuously scans devices to generate updated topology maps and configuration reports.

- **Example:** After a firewall change, all affected network diagrams are automatically updated.

2. Real-Time Anomaly Detection

Compares live traffic against established baselines to flag unauthorized changes.

- **Example:** Detects unusual traffic patterns and immediately initiates Deep Diagnosis.

3. Natural Language Troubleshooting

Processes technical queries in plain English, grounded in your live network state.

- **Example:** "Show me devices with outdated TLS versions" or "Why is VPN throughput dropping at 3 PM daily?"

4. Automated Compliance Auditing

Scans configurations against CIS/NIST benchmarks without human intervention.

- **Example:** Identifies and flags all switches with weak SNMP community strings.

5. ITSM Integration & Ticket-Driven Diagnostics

NetBrain integrates with ITSM platforms to support ticket-driven diagnostics. When a ticket is created:

- Ticket metadata is passed to NetBrain
- Relevant diagnostics and Deep Diagnosis are triggered
- Automation executes immediately
- AI-assisted interpretation summarizes results
- Structured findings and evidence are returned to the originating ticket
- Completion status is recorded

This occurs without human intervention for diagnosis and evidence capture only. Remediation remains governed and requires human approval.

This approach improves L1 and L2 effectiveness while preserving accountability and audit trails.

- **Example:** For a "failed switch port" ticket, NetBrain instantly appends connected device list, last config change, and alternative paths.

6. Accelerated Parser Development

Network engineers spend hours manually building parsers for device outputs. Parser Assistant automates the heavy lifting.

- **Example:** An engineer provides a sample "show version" output; Parser Assistant generates draft parsing logic for review and approval—cutting development time from hours to minutes while maintaining human oversight.

Measurable Business Outcomes

By applying agentic reasoning to targeted workflows, NetBrain helps teams achieve tangible results aligned with Gartner's quantified targets:

- **Reduce human-handled tickets** by 15-25% in areas like wireless triage and drift remediation.
- **Cut MTTR for user-impacting incidents** by 40-60% through explainable, AI-generated diagnostic paths.
- **Decrease L3 escalations** by 20-30% as L1/L2 operators receive agent-generated RCA and remediation plans.
- **Lower ad-hoc troubleshooting time** by 10-15% by shifting first-pass diagnostics into automated test-and-verify plans.

What NetBrain AI Explicitly Does Not Do

As Gartner distinguishes, not all "agents" are agentic. NetBrain deliberately avoids being just a scripted workflow or chatbot assistant. To maintain safety, trust, and auditability:

- Does **not** perform autonomous remediation or execute changes without human approval.
- Does **not** act as a "repackaged assistant"; it is an event-driven planner and executor.
- Does **not** bypass change control or ITSM workflows.
- Does **not** detect cyber threats or replace SIEM/SOAR platforms.
- Does **not** operate on inference alone; every conclusion is grounded in executed diagnostics.
- Does **not** certify or enforce compliance on your behalf.

Conclusion

AI in network operations must reduce execution risk, not amplify it.

NetBrain applies AI where it strengthens consistency, accelerates investigation, and improves evidence quality, while avoiding unbounded autonomy and preserving human accountability.

This is not autonomous networking. It is controlled, operational intelligence.

Gartner, "Innovation Insight: Agentic NetOps Software Bridges Human Intent and Governed Action," Mike Leibovitz, Jonathan Forest, Andrew Lerner, Tim Zimmerman, Nauman Raja, 18 February 2026. (ID G00848904)

Gartner does not endorse any vendor, product or service depicted in its research publications, and does not advise technology users to select only those vendors with the highest ratings or other designation. Gartner research publications consist of the opinions of Gartner's research organization and should not be construed as statements of fact. Gartner disclaims all warranties, expressed or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose.