clarytis

PRODUCT BROCHURE



Argus

VALIDATION PLATFORM

Every AI decision made defensible.

Independent validation, live monitoring and governance for machine learning, generative AI and AI agents, from intake to retirement.

ML, GenAI and AI agents

On-prem, SaaS or air-gapped

RBI · EU AI Act · NIST AI RMF

WHY NOW

AI now makes decisions your organisation has to defend.

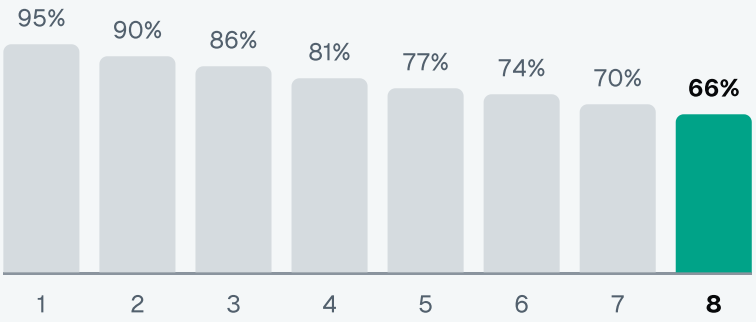
Models approve loans and flag fraud. Generative AI answers customers. Agents act on their own. Each decision needs an owner, evidence that it works and a way to stop it.

COMPOUND ERROR

66%

An agent that gets each step right 95% of the time finishes an eight-step task correctly only about two times in three.

Chance the whole task is still right after each step



Step in an eight-step task. Arithmetic: 0.95 to the power 8 ≈ 0.66.

WHEN AUTOMATED DECISIONS GO WRONG

2021 · PRICING MODEL

Zillow wrote down $304 million of homes in one quarter after its price forecasts misjudged the market.

Source: Zillow Group, Q3 2021 shareholder letter

2024 · CHATBOT

A tribunal held Air Canada liable for refund advice that its website chatbot got wrong.

Source: Moffatt v. Air Canada, 2024 BCCRT 149

2012 · TRADING SYSTEM

Knight Capital lost more than $460 million in 45 minutes when faulty trading software went live.

Source: US SEC, press release 2013-222

REGULATORS NOW EXPECT PROOF

On 24 June 2026 the RBI issued draft Guidance on Regulatory Principles for Model Risk Management. It covers every model a regulated entity uses, including AI/ML and third-party models, at banks, NBFCs and co-operative banks.

Also in view: the EU AI Act, India's DPDP Act, SR 11-7 and NIST AI RMF.

Source: RBI press release, 24 June 2026

MEET ARGUS

One platform to validate, watch and govern every AI system.

Argus is an independent validation and governance layer. It does not build or serve your models. It proves they are safe, fair and correct before they ship, and keeps proving it while they run.

01 · VALIDATE

Before it ships

An independent validator runs a repeatable test battery, and every result is kept as evidence you can re-run.

02 · WATCH

While it serves

The Sentinel runs its own checks on live models. Drift or an expiring certificate triggers a fresh validation.

03 · GOVERN

At every decision

Owners never validate their own models. A council votes by risk tier, and every step lands in a tamper-evident audit trail.

BUILT FOR THREE KINDS OF AI

Classical ML

Performance, calibration, stability, fairness, robustness and explainability.

Generative AI

Faithfulness, hallucination, retrieval quality, safety and prompt-injection resistance.

AI agents

The whole trajectory: coordination, tool and memory use, goal drift, red-teaming and guardrails.

COVERAGE

As of November 2026

22
model families

206
algorithms

177
validation metrics

HOW IT WORKS

Every step from intake to retirement is gated and recorded.

Each model moves through a governed lifecycle. Every gate is held by a named role, and nothing moves without a record.

BEFORE PRODUCTION

01	**Register**	Every model, GenAI app and agent in one inventory, each with a risk tier.
02	**Validate**	An independent validator runs the test battery; the evidence is stored.
03	**Challenge**	AI judges debate and vote; strong dissent blocks automatic sign-off.
04	**Approve**	The validator signs off, the lead approves, the council votes by tier.
05	**Certify**	A certificate whose 'valid until' date follows how fast the model drifts.

IN PRODUCTION

06	**Watch**	The Sentinel runs its own label-free checks on every live model.
07	**Contain**	A critical finding, a tampered model or a lapsed certificate suspends serving at once. People decide the rest.
08	**Retire**	A nine-step governed retirement with a whole-estate impact review.

INDEPENDENCE BY DESIGN

Owners never validate their own models, and never choose who does.

83
permissions, set role by role

9
permanent council seats; only the Chairman holds a veto

0
roles that can edit evidence once it is captured

WHAT SETS ARGUS APART

Eight ways Argus makes an AI decision defensible.

Model Parliament

AI judges from 11 providers, or local models, debate and vote. Dissent is recorded, not averaged away, and above your threshold it blocks automatic sign-off.

Compound-error attribution

Nine agentic engines grade an agent's whole trajectory and point to the step that broke it, not just the final score.

Compliance Passport

One evidence base, mapped to a taxonomy of 19 controls, with an independent verdict for each regulator. 'We cannot tell' is never shown as green.

Temporal certificates

Each certificate's 'valid until' date follows how fast the model drifts, and it escalates automatically if drift crosses its trigger first.

Governed kill switch

Stop one model or the whole AI estate in one action. Restoring it is a council vote, and the Chairman's word is final.

Validation as code

Validation rules live as version-controlled YAML and run in your CI/CD pipeline, so governance has a git history.

Air-gapped by design

Runs fully offline with local AI judges. Every call to an AI model passes an egress guard and is logged, never its text.

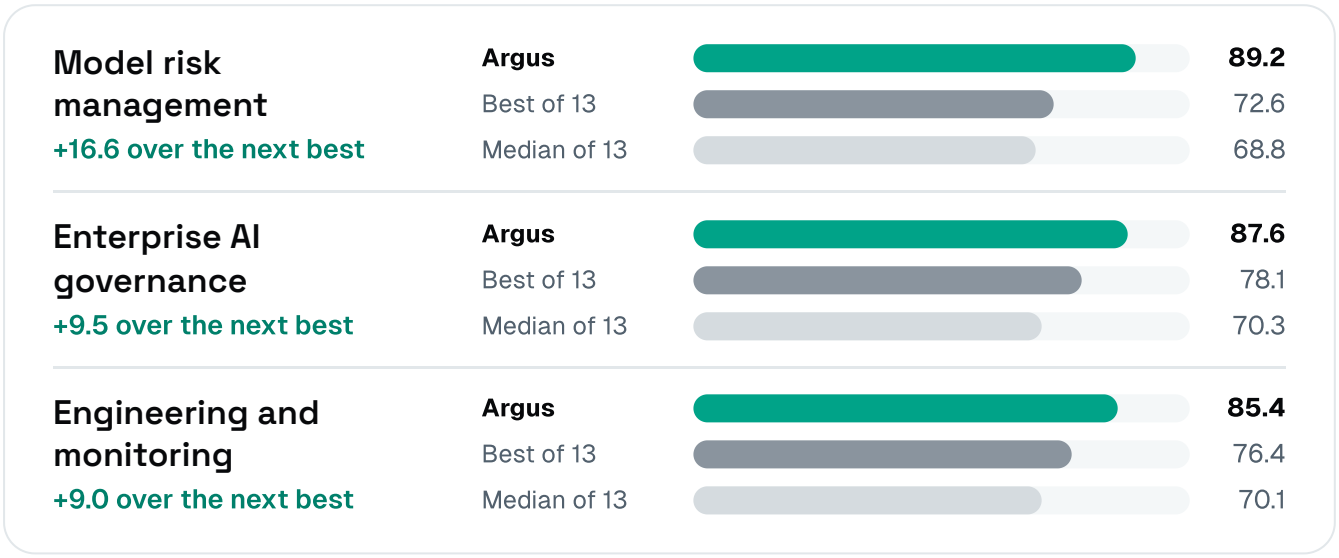
Argus checks itself

Argus re-runs its own platform checks every night and shows the results in plain view. A missed night is flagged, not hidden.

BENCHMARK

Benchmarked against 13 established platforms.

Clarytis scored Argus and 13 platforms on 12 capability categories and 47 feature checks, weighted for three kinds of buyer. Scores are out of 100.



RATED AHEAD OF ALL 13 IN FIVE CATEGORIES

CATEGORY, RATED OUT OF 5	ARGUS	BEST OF 13
Decision independence	●●●●● 5	●●●○○ 3
Model-risk lifecycle	●●●●● 5	●●●●○ 4
Audit evidence and reporting	●●●●● 5	●●●●○ 4
Runtime intervention	●●●●● 5	●●●●○ 4
Customer-controlled deployment	●●●●● 5	●●●●○ 4

Level with the best on four more categories, and rated lower on three: inventory and vendor oversight, policy and compliance, and workflow configuration.

HOW TO READ THIS

Argus is scored on its completed design and named roadmap; the 13 platforms on public evidence dated 14 September 2026 (Arthur AI: 1 October 2026). Scores show fit to requirements, not delivery readiness, security, price or performance. Platforms assessed: Arize AX, Arthur AI, Credo AI, Databricks, Dataiku, Fiddler, Holistic AI, IBM, ModelOp, OneTrust, SAS, ServiceNow and ValidMind.

Source: Clarytis, Argus visual capability benchmark, 1 October 2026

BUILT FOR REGULATED ENTERPRISES

Ready for your regulators, your security team and your auditors.

ONE VALIDATION, EVIDENCE FOR EVERY FRAMEWORK

RBI · SEBI · IRDAI · DPDP Act · EU AI Act · NIST AI RMF · SR 11-7 · ISO/IEC 42001 · GDPR

Each test maps to the clauses it evidences, so one validation serves several regimes. Mappings are versioned and editable, so they follow regulatory change without a software release. Draft rules are tracked but never scored.

DEPLOY WHERE YOUR DATA LIVES

On-premises
In your own data centre.

Private cloud
Your own VPC on Azure or AWS.

Hybrid
Test runners sit beside your sensitive data.

SaaS
Hosted, with strict tenant isolation.

Air-gapped
No outside connection; local AI judges.

One codebase serves every option, with no per-client forks.

CONNECTS TO YOUR STACK

13 model-serving integrations

9 model registries

11 LLM providers for AI judges

18 data sources

SECURITY BUILT IN

- Enterprise SSO (OIDC, SAML) and SCIM
- Sandboxed, network-restricted test runners
- Defences against OWASP and LLM Top 10
- Encryption in transit and at rest
- PII detection and masking
- Append-only, hash-chained audit log

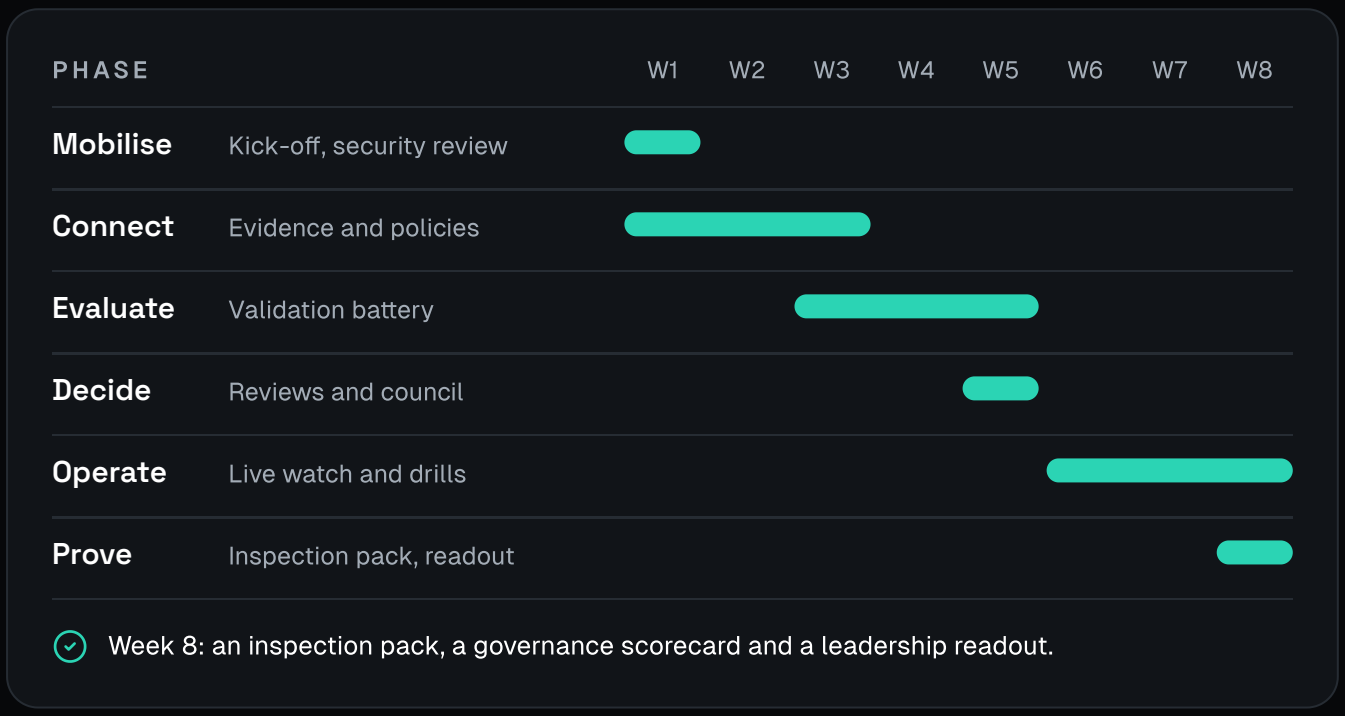
INDUSTRY-AGNOSTIC, WITH ROOTS IN BANKING

Banking and finance · Insurance · Pharma · Automotive · Aviation · Retail · Energy · Telecom · Public sector · Manufacturing

GET STARTED

See Argus govern one of your AI systems in eight weeks.

Start with a proof of value: one live AI system, governed end to end in your own environment. Timings are indicative and agreed at kick-off.



Week 8: an inspection pack, a governance scorecard and a leadership readout.

WHO IT IS FOR

Heads of Model Risk and CROs
One regulator-ready record of every model.

Compliance and internal audit
Framework-mapped reports and certificates.

Boards and executives
Risk posture and a critical-only decision queue.

CISOs
Air-gapped options, a full audit trail and a kill switch.

Chief Data and AI Officers
A governed, visible path from build to production.

Model validators
A full test battery that captures the evidence.

Make every AI decision defensible.

info@clarytis.com · clarytis.com · Whitefield, Bengaluru

Request a demo →

Argus is a Clarytis product. Product figures as of November 2026. © 2026 Clarytis.