

MODEL CARD

BlackSwan Foresight™ Decision-Intelligence Suite

getblack.llc · BlackSwan Foresight, Bengaluru, India

Document type	AI Model Card (governance disclosure)
Product	BlackSwan Foresight™ Suite — 7 models: Maintenance, Quality, SupplyChain, Finance, Revenue, Talent, Customer Foresight™
Version	v0.2 (blackswan-engine-v0.2) — adds Maintenance Foresight™ and Quality Foresight™
Prepared for	getblack.llc — prospect, client, and partner disclosure
Owner	D.V. Shankar, Founder & CEO — BITS Pilani
Governance frameworks referenced	ISO/IEC 42001:2023, NIST AI RMF 1.0
Date	25 July 2026

This model card documents the AI system underlying BlackSwan Foresight™ for prospects, clients, auditors, and referral partners (CAs, auditors) evaluating the product on governance grounds. It follows the structure requested for internal review and is written to be publishable on getblack.llc or attached to a client due-diligence packet without disclosing proprietary implementation detail. Version 0.2 adds the two new shop-floor modules, Maintenance Foresight™ and Quality Foresight™, and makes explicit that these two — unlike the original five — require a data source beyond Tally.

1. Model Purpose

BlackSwan Foresight™ is a decision-intelligence system that converts a manufacturer's existing Tally accounting data — and, for two of its seven models, existing shop-floor sensor or camera data — into forward-looking forecasts and recommended actions, without requiring an ERP migration or an in-house data-science function. It is aimed at promoter-owned Indian manufacturers in the ₹40–80 crore revenue band.

The suite is not a single model but seven purpose-built Foresight™ models, each addressing a distinct promoter decision:

Model	Intended use	Primary decision supported
Maintenance Foresight™ (flagship, new)	Predictive maintenance — breakdown prediction, maintenance-schedule and spare-parts recommendations from Tally + existing PLC/IoT sensor data	When a machine is likely to fail, and what maintenance/parts action to take first
Quality Foresight™ (flagship, new)	AI vision inspection for surface, assembly, and packaging defects, linked to Tally inventory/quality records	Which batches or lines have a rising defect rate, and where to intervene
SupplyChain Foresight™	Demand forecasting, inventory optimization, supplier-risk flagging	What to stock, what to stop stocking, which supplier is at risk
Finance Foresight™	Cash-flow shortfall prediction, scenario planning, financial risk scoring	When a cash gap will occur and what to do before it does
Revenue Foresight™	Sales forecasting, market-trend signal detection, customer lifetime value	Where the order book is heading and which accounts to prioritize
Talent Foresight™	Workforce planning, attrition prediction, skill-gap analysis	Where the floor is understaffed or about to lose people
Customer Foresight™	Churn prediction and segmentation	Which accounts are quietly slipping while there is still time to act

Out of scope: BlackSwan Foresight™ is not a general-purpose ERP, an accounting system of record, a credit-underwriting or lending decision tool, a certified safety-instrumented system, or a replacement for statutory audit or

compliance filings. Maintenance Foresight™ does not replace certified safety interlocks or statutory equipment inspection, and Quality Foresight™ does not replace certified quality-control instrumentation where regulation requires it. It does not make autonomous financial or operational commitments on the client's behalf — every output is a forecast or recommendation for a human decision-maker (the promoter or their finance/ops/plant lead).

2. Training Data

2.1 Per-client input data

Each deployment is trained/calibrated on the client's own historical Tally ledger export (sales, purchase, inventory, and receivables/payables registers), typically 24–36 months of transaction history where available [confirm exact minimum window used in production]. For Maintenance Foresight™ and Quality Foresight™, this Tally history is combined with a second, client-specific data source that does not live in Tally: machine/sensor telemetry from the client's existing PLCs or IoT sensors (Maintenance), or camera/image captures at an inspection point (Quality). All data is client-specific and is not pooled across clients or used to train any shared or public model — client data segregation is a stated commitment on getblack.llc's FAQ and privacy language, and now explicitly extends to this shop-floor data.

2.2 Nature and volume

Attribute	Detail
Source (5 original models)	Client-exported Tally ledgers (CSV/XLS export from Tally Prime/ERP 9)
Source (Maintenance Foresight™)	Existing PLC/IoT sensor feed (e.g. vibration, temperature, runtime-hours) — not captured in Tally — plus Tally purchase/maintenance-job records
Source (Quality Foresight™)	Camera or IP-camera image captures at an inspection point — not captured in Tally — plus Tally inventory/quality records
Data types	Sales register, purchase register, stock/inventory ledger, AR/AP ageing, (optionally) payroll/HR register for Talent Foresight™; sensor telemetry stream for Maintenance; image stream for Quality
Volume per client	Typically 10k–100k+ line-item transactions across the trailing 24–36 months for the Tally-based models [confirm — varies by client size]; sensor/image volume for Maintenance and Quality varies with sensor count and line throughput and has not yet been standardized [confirm minimums]
Granularity	SKU-level for SupplyChain/Revenue; ledger-entry level for Finance; employee-record level for Talent; machine/asset-level for Maintenance; batch or frame-level for Quality
PII / sensitive fields	Employee names/compensation (Talent), customer names (Revenue/Customer) — handled as sensitive per client data-handling policy. Camera captures for Quality Foresight™ may incidentally include personnel in frame and are treated as sensitive by the same policy
Synthetic/demo data	A synthetic demo dataset (illustrative footwear manufacturer, ~₹55 Cr trailing revenue) is used for sales demonstrations only — clearly separated from production client data. No production Maintenance/Quality sensor or image dataset exists yet [confirm demo-data plan for the two new modules]

2.3 What the model does NOT train on

- No client data — including Tally data, sensor telemetry, or camera imagery — is used to train, fine-tune, or improve any shared/public AI model; each client's forecasts are generated from that client's own data only.
- No data is used to train third-party foundation models used within the architecture (see Section 3) beyond standard, non-retained inference calls.
- Statistical/ML components (see 3.2) are fit per-client at onboarding and refreshed on a defined cadence [confirm refresh interval — e.g. monthly] rather than trained once on a static global corpus.

3. Model Architecture

BlackSwan Foresight™ is built as an agentic, MCP-based (Model Context Protocol) system — a coordinated set of specialized agents and analytical modules, rather than one monolithic model. This is described on getblack.llc as "not a dashboard — a team of agents."

3.1 System layers

Layer	Function
Ingest	Parses and normalizes Tally exports (sales, purchase, inventory, AR/AP, HR) into a common schema; for Maintenance/Quality Foresight™, also ingests PLC/IoT sensor telemetry or camera image captures from a separate feed
Diagnostics	Data-quality checks, anomaly detection, and readiness scoring (feeds the AI-readiness assessment)
Forecasting	Time-series and ML forecasting per model — demand, cash flow, revenue, attrition, churn, machine breakdown
Scoring	Risk/priority scoring — supplier risk, financial risk, churn risk, attrition risk, breakdown risk, defect rate
Reporting	Generates the client-facing read-out: forecast, flagged exceptions, and a recommended next action

3.2 Analytical / ML techniques

Forecasting and scoring draw on a mix of classical statistical and applied ML methods appropriate to tabular, ledger-style data — e.g. ARIMA-family time-series models for demand and cash-flow forecasting, tree-based methods (e.g. XGBoost/CART) for risk scoring and churn/attrition classification, and nearest-neighbour/segmentation methods for customer and SKU clustering. Maintenance Foresight™ adds anomaly-detection and survival-analysis style methods over sensor telemetry to estimate time-to-failure. Quality Foresight™ adds an edge computer-vision (CNN-based) defect-classification model over camera captures [confirm specific model/vendor and whether inference runs on-device or is sent to a cloud endpoint — this materially affects the data-handling disclosure in 5.5].

3.3 Agentic / orchestration layer

A seven-tool MCP server exposes the ingest, diagnostics, forecasting, scoring, and reporting capabilities as callable tools. An orchestrating agent (deployed via Claude Desktop / Claude for enterprise use in the reference implementation) sequences these tools per client request — e.g. "run a SupplyChain read-out" triggers ingest → diagnostics → forecast → score → report in sequence, with the agent selecting which tool to call based on the question asked. This is the "team of agents" referenced on the product site. [Confirm whether the MCP tool count/interfaces have been extended to cover Maintenance/Quality ingest, or whether this remains a roadmap item.]

3.4 Design considerations

- No new ERP required — the architecture is deliberately built to read Tally exports rather than requiring system integration, which is the basis of the ~2-week time-to-value claim for the five Tally-only models.
- Maintenance and Quality Foresight™ require a secondary data source beyond Tally — an existing PLC/IoT sensor feed, or a mobile/IP camera at an inspection point, respectively. Clients without that infrastructure can still adopt the other five Tally-only models; the ~2-week time-to-value claim does not automatically extend to Maintenance/Quality until that infrastructure is confirmed present.
- Per-client isolation — each client's models and data are kept segregated rather than pooled, both for confidentiality and because promoter-owned manufacturers are highly heterogeneous (no single global model would generalize well across them).
- Human-in-the-loop by design — outputs are forecasts and recommendations, not automated actions; a promoter or finance/ops/plant lead reviews and acts.
- Governance alignment — architecture and data handling are informed by ISO/IEC 42001 and the NIST AI RMF, reflecting the founder's AIGP certification, rather than being built and governed as an afterthought.

4. Inputs and Outputs

4.1 Inputs

Model	Required input	Format
Maintenance Foresight™	PLC/IoT sensor feed (e.g. vibration, temperature, runtime-hours) + Tally purchase/maintenance-job records	Sensor telemetry stream/CSV + Tally CSV/XLS export
Quality Foresight™	Camera or IP-camera image captures at an inspection point + Tally inventory/quality records	Image stream (JPEG/PNG) + Tally CSV/XLS export
SupplyChain Foresight™	Sales register, purchase register, stock/inventory ledger	Tally CSV/XLS export, SKU-level
Finance Foresight™	AR/AP ageing, bank/cash ledger, sales & purchase registers	Tally CSV/XLS export, ledger-entry level
Revenue Foresight™	Sales register, customer master, order history	Tally CSV/XLS export, transaction-level
Talent Foresight™	Payroll/HR register, headcount and attendance data (optional module)	CSV/XLS, employee-record level
Customer Foresight™	Sales register, customer master, invoice/payment history	Tally CSV/XLS export, customer-level

4.2 Outputs

- Forecasts: e.g. monthly demand by SKU (₹ and units), projected cash position by week/month, revenue/order-book projection, predicted machine breakdown window.
- Risk/priority scores: supplier-risk score, financial-risk score, churn-risk score, attrition-risk score, breakdown-risk score, defect-rate score — typically normalized 0–100 [confirm scale] with a plain-language risk band (low/medium/high).
- Flagged exceptions: e.g. named slow-moving SKUs and estimated overstock value, named accounts at churn risk, named suppliers flagged for risk, named machines approaching predicted failure, batches with a rising defect rate.
- Recommended next action: a short, specific recommendation attached to each flagged item — for Maintenance Foresight™, this includes an automatically generated work-order or spare-parts purchase suggestion written back into Tally as a draft transaction, not an automatic order.
- Read-out format: web-based report / assessment result in the client-facing product; underlying data returned as structured JSON from the MCP tool layer for downstream use.

4.3 Expected ranges / assumptions to disclose

Forecast outputs are probabilistic projections, not guarantees, and accuracy depends on the completeness and consistency of the client's Tally data. For Maintenance Foresight™ and Quality Foresight™, forecast and detection quality additionally depends on the coverage, placement, and data quality of the client's existing sensors or camera — this is a second point of failure beyond Tally data quality and should be scoped explicitly during discovery, not assumed.

5. Integration Points

5.1 Data feeds

Two distinct data-feed types now exist. (a) Client Tally exports — the primary and, for five of the seven models, sole data source (manual export/upload at onboarding and on a refresh cadence [confirm: manual re-upload vs. scheduled export]; no live/API connection into Tally is implied by current site language). (b) Shop-floor feeds — existing PLC/IoT sensor data (Maintenance Foresight™) or camera/image captures (Quality Foresight™) — required in addition to Tally for those two models specifically. Clients without existing PLC/IoT or camera

infrastructure are not currently supported for Maintenance or Quality Foresight™ on a Tally-only basis; this should be stated plainly to prospects rather than implied to work "out of the box" the way the other five models do.

5.2 APIs and tooling

- MCP server (seven tools) exposing ingest, diagnostics, forecasting, scoring, and reporting functions — the integration surface for the orchestrating agent and for any future direct API consumption. [Confirm whether sensor/camera ingest tools have shipped or remain roadmap.]
- Reference orchestration client: Claude Desktop / Claude, used in the current implementation to run the agent that calls the MCP tools.

5.3 Client-facing touchpoints

- AI-readiness assessment (web form on getblack.llc): questionnaire that produces a readiness score and a recommended first Foresight™ model, now including Maintenance and Quality as possible recommendations — routes to email and a discovery-call booking.
- Discovery call → data handoff → onboarding: the manual handoff points where a prospect's raw data enters the system are today largely human-mediated (email/file transfer for Tally exports), not an automated pipeline [confirm current state, and confirm the handoff process for sensor/camera data specifically, which has not been defined in prior product materials — flag as a manual/undefined process for governance purposes].

5.4 Internal/organizational systems

- Outbound and CRM stack — upstream of the model, used for lead generation and does not feed the forecasting models.
- No current integration with client ERP/finance systems beyond Tally export ingestion; no SAP/Oracle connector exists today, consistent with the "0 new ERP" positioning. No current integration with a specific PLC/IoT platform or camera vendor exists today either — Maintenance and Quality Foresight™ should be scoped per client's existing hardware until a standard integration is built [confirm].

5.5 Data handling boundary

Per the published FAQ: client data is not used to train any public AI model, financial/workforce data is kept segregated per client, and data handling follows a documented approach informed by NIST AI RMF and ISO 42001. This boundary now explicitly extends to shop-floor data: sensor telemetry and camera imagery are treated with the same segregation commitment as financial and HR data, and camera imagery in particular should be governed by a stated retention/deletion policy given it may incidentally capture personnel [confirm retention period and whether raw images are stored or only defect classifications]. Any model card published externally should keep this boundary explicit, since it is the primary trust claim being made to promoters handling sensitive financial, HR, and now shop-floor data.

6. Limitations, Risks & Review

- Forecast quality is bounded by Tally data quality/completeness at each client — not independently verified against ground truth at scale yet [confirm status of any pilot accuracy validation].
- Maintenance Foresight™ and Quality Foresight™ depend on the client having, or installing, compatible PLC/IoT sensors or a camera respectively. These are prerequisites, not currently bundled hardware, and should be scoped during discovery before any accuracy or time-to-value commitment is made — this is a materially different sales motion from the Tally-only models.
- Small-sample risk: manufacturers with limited transaction history or highly seasonal/one-off demand patterns may get lower-confidence forecasts — recommend disclosing a minimum-data threshold once defined. The same applies to Maintenance/Quality with limited sensor history or inspection volume.

- Neither Maintenance Foresight™ nor Quality Foresight™ is a certified safety or quality-control instrument; outputs are decision support and should not be relied upon in place of statutory safety or quality-compliance systems.
- Human oversight required: outputs are decision support, not autonomous decisions — this should remain a standing product principle, not just a disclosure line.