MODEL CARD

BlackSwan Foresight™ Decision-Intelligence Suite

blackswanforesight.in · BlackSwan Foresight, Bengaluru, India

FIELD	VALUE
Document type	AI Model Card (governance disclosure)
Product	BlackSwan Foresight™ Suite: 4 models (SupplyChain, Finance, Revenue, Customer Foresight™) plus Prebuilt Scenario Forecasts
Version	v0.5 (blackswan-engine-v0.6; scenario calculator v1). Adds Prebuilt Scenario Forecasts (Section 7) and the Data Trust diagnostics description; moves the site to blackswanforesight.in. The four Foresight™ models and their methods are unchanged since v0.3
Prepared for	blackswanforesight.in: prospect, client and partner disclosure
Owner	D.V. Shankar, Founder & CEO, BITS Pilani
Contact	shankar.dv@blackswanconsulting.in
Governance frameworks referenced	ISO/IEC 42001:2023, NIST AI RMF 1.0 (used as design guidance; not a certification claim)
Date	26 September 2026

This document describes the AI system underlying BlackSwan Foresight™ for prospects, clients, auditors, and referral partners (CAs, auditors) evaluating the product on governance grounds. Where a claim describes the intended production design rather than something already verified against real client data, that distinction is stated explicitly rather than implied.

1. Model Purpose

BlackSwan Foresight™ is a decision-intelligence system that converts a manufacturer's existing Tally or ERP data into forward-looking forecasts and recommended actions, without requiring an ERP migration or an in-house data-science function. It is aimed at mid-market Indian manufacturers (roughly ₹40–500 crore revenue), typically promoter-led.

The suite is four purpose-built Foresight™ models, each addressing a distinct decision, plus a scenario layer that stress-tests liquidity:

MODEL / LAYER	INTENDED USE	PRIMARY DECISION SUPPORTED
SupplyChain Foresight™	Demand forecasting, inventory reorder optimization, supplier-risk flagging	What to stock, what to stop stocking, which supplier is at risk
Finance Foresight™	Working-capital teardown, cash-flow risk, ledger anomaly detection	When a cash gap will occur, and which transactions warrant review

MODEL / LAYER	INTENDED USE	PRIMARY DECISION SUPPORTED
Revenue Foresight™	Sales forecasting, cash-conversion tracking, customer default-risk scoring	Where the order book is heading and which accounts to prioritize collecting from
Customer Foresight™	Concentration-risk analysis, customer default-risk scoring	Which accounts are a concentration or default risk while there's still time to act
Prebuilt Scenario Forecasts	Deterministic stress testing of liquidity against eight named preshocks (Section 7)	How much pressure the business can take before a cash gap opens, and what to do first

Out of scope: BlackSwan Foresight™ is not a general-purpose ERP, an accounting system of record, a credit-underwriting or lending decision tool, or a replacement for statutory audit or compliance filings. It does not make autonomous financial commitments on the client's behalf — every output is a forecast or recommendation for a human decision-maker (the promoter or their finance/operations lead).

2. Training Data

2.1 Per-client input data (production design)

Each deployment is intended to be 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]. Exports from other ERPs (CSV/XLS) are intended to be mapped to the same schema during onboarding [confirm — not yet tested on non-Tally data]. Ledger-derived forecasts (SupplyChain, Finance) are fit per client and are not pooled across clients.

One exception, disclosed explicitly: the credit-risk model shared by Revenue Foresight™ and Customer Foresight™ (see 3.2) is a single trained model, not refit per client. In its current form it is trained on a synthetic population, not pooled real client data (see 2.4) — no client's receivables data is used to train another client's score. [confirm production retraining/pooling policy once real client payment-outcome data exists]

Prebuilt Scenario Forecasts are not trained on any data. They are deterministic arithmetic applied to the inputs described in 4.1 and Section 7.

2.2 Nature and volume

ATTRIBUTE	DETAIL
Source	Client-exported Tally ledgers (CSV/XLS export from Tally Prime/ERP 9); other ERP exports mapped at onboarding [confirm]
Data types	Sales register, purchase register, stock/inventory ledger, SKU usage history, AR/AP ageing
Volume per client	Typically 10k–100k+ line-item transactions across the trailing 24–36 months [confirm — varies by client size; larger mid-market clients may exceed this range]
Granularity	SKU-level (SupplyChain); ledger-entry/transaction-level (Finance); customer-level (Revenue, Customer)
PII / sensitive fields	Customer names (Revenue/Customer Foresight) — handled as sensitive per client data-handling policy

2.3 What the model does NOT train on

- No client data — including Tally exports — is used to train, fine-tune, or improve any shared/public AI model; each client's ledger-derived forecasts (SupplyChain, Finance) are generated from that client's own data only.
- No data is used to train third-party foundation models used within the architecture beyond standard, non-retained inference calls.
- SupplyChain and Finance's statistical components are fit per client at onboarding and refreshed on a defined cadence [confirm refresh interval] rather than trained once on a static global corpus. The Revenue/Customer credit-risk model is the exception — see 2.1 and 2.4.
- Numbers entered into the Prebuilt Scenario Forecasts calculator on the website are processed only in the visitor's browser; they are not transmitted to or stored by BlackSwan Foresight.

2.4 Current implementation status — synthetic data only

As of this version (blackswan-engine-v0.6), the engine has been built and tested exclusively against synthetic, algorithmically generated mock data — no live Tally ODBC connection and no real client data have been processed by the system yet. This applies across all four models:

- SupplyChain Foresight™'s demand forecasts and safety-stock calculations run on synthetic monthly usage histories, seeded per demo tenant.
- Finance Foresight™'s anomaly detection runs on a synthetic transaction set with a small number of deliberately injected outliers, so the detector has something concrete to find.
- The Revenue/Customer credit-risk model is trained on a synthetic labeled population (`synthetic_data.credit_training_population()`) with a logistic ground-truth relationship between features and default, not real invoice outcomes. Held-out AUC on this synthetic set is approximately 0.69 — meaningfully above a random baseline, confirming the model is learning the intended feature relationship, but this figure describes fit quality on synthetic data, not real-world predictive accuracy.
- Prebuilt Scenario Forecasts run today as a published calculator on blackswanforesight.in, on numbers the user enters (prefilled with an illustrative ₹60 crore manufacturer). Feeding the scenario layer automatically from a client's Tally data and sanction terms through the MCP engine is the production design and has not yet been built.

This is a material fact for governance evaluation and is stated plainly rather than implied: the architecture and techniques described in this document are real and tested, but production use on a live client's actual Tally data has not yet occurred. Section 5.1 covers the live-ODBC integration gap; Section 6 covers the resulting limitations.

3. Model Architecture

BlackSwan Foresight™ is built as an agentic, MCP-based (Model Context Protocol) system — a coordinated set of specialized tools, rather than one monolithic model. This is described on blackswanforesight.in 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) into a common schema

LAYER	FUNCTION
Diagnostics (Data Trust)	Data-quality checks (profiling for gaps, duplicates and outliers; validation against accounting rules) and transaction-level anomaly detection (isolation forest, Finance Foresight™). Reconciliation against GST returns and bank statements, and a per-dataset trust score, are part of the published Data Trust process [confirm whether performed automatically or by an analyst in the current build]
Forecasting	Time-series and statistical forecasting per model — demand (Holt-Winters), inventory reorder points (safety-stock formula, with TSB for intermittent-demand SKUs)
Scoring	Risk/priority scoring — stockout/overstock risk, transaction anomaly score, customer default-risk score
Scenario	Prebuilt Scenario Forecasts: deterministic liquidity projection under eight named preshocks, at three severities (Section 7)
Reporting	Generates the client-facing read-out: forecast, flagged exceptions, scenario readout and a recommended next action

3.2 Analytical / statistical / ML techniques

Each model uses a technique matched to its own data pattern, rather than one generic method applied uniformly:

SupplyChain Foresight™ — demand is forecast with Holt-Winters triple exponential smoothing (level + trend + 12-month seasonal component), fit on rolling monthly sales history and validated with a held-out backtest. Backtested forecast error (MAPE) has been in the 5–13% range in testing on synthetic data with sufficient history [confirm on real client data]. Reorder points use a statistical safety-stock calculation (safety stock = $z \cdot \sigma(\text{daily usage}) \cdot \sqrt{(\text{lead-time days})}$) rather than a fixed threshold. SKUs are additionally classified by demand pattern using the Syntetos-Boylan framework (ADI/CV^2 thresholds, the field-standard cutoffs of 1.32/0.49); items with genuinely intermittent or lumpy demand are forecast with TSB (Teunter-Syntetos-Babai) smoothing instead of the mean/variance formula, since the normal-distribution assumption behind standard safety stock does not hold for spare-parts-style start-stop demand. SKUs are further classified by ABC-XYZ (value × predictability, Pareto cutoffs) to set a differentiated target service level (z-score) per item rather than one constant across the catalog. ABC ranking is a small-catalog caveat: with fewer than roughly 20–30 SKUs, Pareto cumulative-share cutoffs can place an individually high-value SKU in a lower class than intuition suggests — mathematically correct, but worth re-validating once a real client's full catalog is available.

Finance Foresight™ — the working-capital teardown (trapped cash, working-capital loan coverage) is deterministic arithmetic on accounting identities, not a model — kept that way deliberately, since auditable, explainable arithmetic is preferable to a model for this calculation. Transaction-level anomaly detection uses Isolation Forest, an unsupervised technique (scikit-learn, contamination parameter 0.03 as a starting estimate [confirm once real transaction volumes are known]), chosen because no labeled fraudulent- or erroneous-transaction data exists to train a supervised classifier on.

Revenue Foresight™ and Customer Foresight™ — share one trained Logistic Regression model that outputs a probability of default per customer from payment-history features (outstanding balance, historical late-payment rate, historical average days overdue, and portfolio concentration share), replacing a static days-overdue cutoff. See 2.4 for the current synthetic-training caveat.

Prebuilt Scenario Forecasts — deterministic liquidity arithmetic, deliberately not a learned model, for the same auditability reason as the working-capital teardown. See Section 7.

3.3 Agentic / orchestration layer

A five-tool MCP server (blackswan-engine-v0.6) exposes the four Foresight models plus a connector-status utility tool as callable tools. An orchestrating agent sequences these tools per client request — e.g. “run a SupplyChain read-out” triggers ingest → forecast → score → report in sequence. The current implementation runs over local stdio transport, suitable for development and demonstration; production multi-tenant HTTP/SSE deployment with per-tenant authentication is designed but not yet built [status: open engineering item, see Section 5.2]. The scenario layer is not yet exposed as an MCP tool.

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 design target once the live-ODBC gap in Section 5.1 is closed.
- **Method-per-problem, not one-size-fits-all**: each model uses the technique suited to its own data pattern (time-series smoothing, safety-stock arithmetic, unsupervised anomaly detection, supervised classification, or deterministic scenario arithmetic) rather than a single generic algorithm.
- **Per-client isolation for ledger-derived models** — SupplyChain and Finance’s statistical components are fit per client; the shared credit-risk model is the disclosed exception (see 2.1).
- **Human-in-the-loop by design** — outputs are forecasts and recommendations, not automated actions; a promoter or finance/operations lead reviews and acts.
- **Governance alignment** — architecture and data handling are informed by ISO/IEC 42001 and the NIST AI RMF, and by the founder’s AI Security & Governance certification (Securiti), rather than being built and governed as an afterthought. The company does not claim certification against ISO/IEC 42001 or the NIST AI RMF.

4. Inputs and Outputs

4.1 Inputs

MODEL	REQUIRED INPUT	FORMAT
SupplyChain Foresight™	Sales register, purchase register, stock/inventory ledger, SKU usage history	Tally CSV/XLS export, SKU-level
Finance Foresight™	AR/AP ageing, bank/cash ledger, sales & purchase registers, transaction-level ledger detail	Tally CSV/XLS export, transaction level
Revenue Foresight™	Sales register, customer master, order history, receivables ageing	Tally CSV/XLS export, transaction-level
Customer Foresight™	Sales register, customer master, invoice/payment history, receivables ageing	Tally CSV/XLS export, customer-level
Prebuilt Scenario Forecasts	17 business parameters (revenue, margins, working-capital days, supplier and customer mix, cash, CC/OD limit and utilisation, debt service, DSCR covenant)	On the website: entered by the user. Production design: derived from the client’s Tally data and bank sanction terms

4.2 Outputs

- **Forecasts:** monthly demand by SKU (units and ₹), dynamic reorder point and safety-stock level per SKU, projected cash position, revenue/order-book projection.

- **Risk/priority scores:** stockout/overstock risk flag per SKU, transaction anomaly score (isolation forest), customer probability of default (0–100%, with a low/watch/high band).
- **Flagged exceptions:** named slow-moving or at-risk-of-stockout SKUs, named anomalous ledger transactions, named accounts at elevated default risk, receivables concentration risk.
- **Scenario readout:** month in which a cash gap opens, runway in months, DSCR against covenant, liquidity headroom at month 12, an 18-month liquidity projection against baseline, and (at Black Swan severity) the breaking point.
- **Recommended next action:** a short, specific recommendation attached to each flagged item and each scenario.
- **Read-out format:** structured JSON returned from the MCP tool layer; the reference web dashboard renders this as a client-facing report. The scenario calculator renders directly in the browser.

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.

SupplyChain Foresight™'s seasonal demand model specifically requires more than 24 months of monthly sales history to fit a 12-month seasonal component at all, and meaningfully more (on the order of 30+ months) to validate it with a rolling backtest — this was confirmed by testing against a real third-party retail dataset with 25 months of history, where the seasonal fit failed in 5 of 6 rolling backtest windows. Clients with shorter sales history do not currently receive a graceful fallback (e.g. a trend-only model without the seasonal component) — this is an open engineering item, not yet implemented as of blackswan-engine-v0.6 — and should be scoped explicitly during discovery rather than assumed to work.

Scenario readouts are only as good as their inputs and the stated assumptions in Section 7.3; they are sensitivity analyses, not forecasts of what will happen.

5. Integration Points

5.1 Data feeds

Client Tally exports are the primary and current sole data source across all four models. As of this version, no live/API (ODBC) connection into Tally exists in production — the engine's adapter layer has a defined interface for live Tally ODBC queries, but the query methods themselves are unimplemented placeholders in the current build. Production use today would rely on a manual export/upload workflow once that path is built, or continue running in the same synthetic-data mode used for demos. [confirm target date for live-ODBC implementation and manual-export fallback design]

5.2 APIs and tooling

- MCP server (five tools: four Foresight models + connector-status) — the integration surface for the orchestrating agent and for any future direct API consumption.
- Reference orchestration client: Claude Desktop / Claude, used in the current implementation to run the agent that calls the MCP tools.
- Transport: local stdio in the current build; production HTTP/SSE with per-tenant authentication is designed but not yet implemented (see 3.3).

5.3 Client-facing touchpoints

- **AI-readiness assessment** (web form on blackswanforesight.in): questionnaire that produces a readiness score and a recommended first Foresight™ model from the four-model suite — routes to email and a discovery-call booking.

- **Prebuilt Scenario Forecasts calculator** (blackswanforesight.in/scenario-forecasts/): runs entirely in the visitor's browser; inputs are not transmitted or stored.
- **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].

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. Clients on ERPs other than Tally would currently need a manual CSV/XLS export mapped by BlackSwan at onboarding [confirm].

5.5 Data handling boundary

Per the published FAQ: client data is not used to train any public AI model, and financial and customer data is kept segregated per client, following a documented approach informed by NIST AI RMF and ISO 42001. This is the production-design commitment; as noted in 2.4, no live client data has been processed by the system as of this version, so the segregation commitment has not yet been exercised against real client data in production.

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, since no production client data has been processed as of this version (see 2.4) [confirm status of any pilot accuracy validation once live].
- The Revenue/Customer credit-risk model is trained exclusively on a synthetic labeled population, not real historical payment outcomes (see 2.4, 3.2). Its probability-of-default output should be treated as illustrative of the scoring approach, not a calibrated real-world risk estimate, until retrained on genuine client payment history via `scripts/train_credit_risk_model.py`.
- SupplyChain Foresight™'s seasonal demand model requires more monthly sales history than some clients may have (see 4.3); shorter-history clients may currently receive a failed model fit or an unreliable seasonal forecast rather than a graceful fallback — an open engineering item, not yet resolved as of blackswan-engine-v0.6.
- No live Tally ODBC integration exists yet (see 5.1) — the ~2-week time-to-value claim describes the intended design once that integration and a defined data-handoff workflow are built, not a currently-verified client onboarding time.
- Scenario readouts rely on simplifying assumptions (Section 7.3). They test one preshock at a time, and the website calculator uses the inputs the visitor enters, not verified books.
- 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.
- ABC classification is a small-catalog caveat (see 3.2) — re-validate behavior once a real client's full SKU catalog is available.
- Mid-market scale: the engine has not yet been exercised on real data at any scale, including the higher transaction volumes and multi-entity or multi-plant structures typical of larger mid-market manufacturers. Multi-entity consolidation behavior should be scoped during discovery [confirm].

- Human oversight required: outputs are decision support, not autonomous decisions — this should remain a standing product principle, not just a disclosure line.

7. Prebuilt Scenario Forecasts NEW IN v0.5

7.1 Purpose and method

Prebuilt Scenario Forecasts measure how much pressure a manufacturer's liquidity can absorb, rather than predicting whether a shock will happen. Each of eight named preshocks is applied at three severities: **Bad** and **Worse** use preset intensities; **Black Swan** is a reverse stress test that searches (by bisection) for the smallest intensity that opens a cash gap within 12 months.

Liquidity is defined as cash plus undrawn CC/OD, where the usable limit is the lower of the sanctioned limit and drawing power (75% of inventory + 60% of receivables), minus the amount already drawn; any amount drawn above that ceiling must be repaid from cash. The calculator projects liquidity month by month for 24 months. A preshock acts through three channels: one-time working-capital absorption and one-time profit loss, both phased in over one to six months; recurring monthly margin loss; and changes to drawing power or the sanctioned limit. Every formula is shown to the user, with the actual numbers substituted, under "How this is calculated".

7.2 The eight preshocks

PRESHOCK	PARAMETER	BAD / WORSE PRESETS	CALCULATION
Collections Drift	Days added to DSO	+15 / +30 days	Cash absorbed = daily revenue × extra days, over 3 months; overdue receivables excluded from drawing power
Big Buyer Stretch	Days the largest customer pays late	+60 / +90 days	Daily revenue × top-customer share × extra days, over 2 months
Festive Demand Miss	% of festive-season sales missed	15% / 30%	Lost sales × contribution margin (profit); unsold stock at material cost (working capital), 75% of which remains eligible for drawing power
GST / ITC Blockage	% of monthly input tax credit blocked for 6 months	20% / 40%	Monthly ITC (purchases × 18% ÷ 12) × share blocked × 6 months; a timing effect, not a loss
45-Day Squeeze	% of purchases that must be paid within 45 days (Section 43B(h))	MSME share / MSME share + 30 pts	Daily purchases × share × (supplier days − 45), over 2 months
Limit Cut	% cut in sanctioned CC/OD limit	15% / 30%	Immediate; if the new limit is below the amount drawn, the excess is repaid from cash
RM Spike	% rise in raw-material prices	10% / 20%	Monthly margin loss = material cost × rise × (1 − pass-through) ÷ 12; stock revaluation effect on working capital and drawing power
FX on Imported RM	% depreciation of the rupee	5% / 10%	As RM Spike, applied to the imported share of material cost

7.3 Assumptions

- The baseline assumes business-as-usual operating cash flows net to zero, so the readout isolates the effect of the preshock.
- Drawing-power margins are fixed at 25% on stock and 40% on receivables; actual sanction terms vary by bank and client.
- GST input credit is taken at 18% of purchases; overdue receivables are treated as ineligible for drawing power.
- Interest on additional borrowing, tax effects and second-order responses (for example, customers or suppliers reacting to one another) are not modelled.
- DSCR = (EBITDA – profit lost within 12 months) ÷ annual debt service, compared with the covenant the user enters.

7.4 Status

> Implemented and published as a browser-based calculator (scenario calculator v1) on blackswanforesight.in/scenario-forecasts/, running on user-entered inputs. Automatic population of the inputs from a client's Tally data and sanction letter, and exposure as an MCP tool, are the production design and are not yet built. Scenario outputs are decision support, not lending, credit or investment advice.

Questions, governance requests, or corrections: shankar.dv@blackswanconsulting.in