

CreditVantage AI

Better Credit Insights. Stronger Lending Decisions.

Technical Methodology & Governance Paper

An institutional overview of our credit-risk methodology, analytical controls and governance framework, and the defined role of artificial intelligence within the AI Credit Analyst.

Version 1.0

September 2026

Document Status and Version

Document Title	AI Credit Analyst Technical Methodology & Governance Paper
Document Owner	CreditVantage AI
Document Classification	Public Document
Document Status	Final
Publication Status	Initial Public Release
Document Version	1.0 Final
Publication Date	20 September 2026
Methodology Scope	Institutional consumer and retail credit-risk decision support
System / Testing Reference	Production methodology and implementation baseline current as at 17 September 2026

Document-Control Statement

This document describes the AI Credit Analyst methodology, governance principles, analytical controls and defined use of artificial intelligence as at the stated publication date. CreditVantage AI may revise its methodology, controls, implementation or governance arrangements over time. Material changes affecting the substance of this paper will be reflected through appropriate document-version updates.

Disclosure Classification

This paper provides a high-level explanation of the AI Credit Analyst methodology and governance framework. It does not disclose detailed proprietary decision rules, calibration, thresholds, implementation logic, internal schemas, confidential test cases or other commercially sensitive methodology.

Table of Contents

1. Executive Overview	1
2. How the AI Credit Analyst Is Designed	2
3. The Credit-Risk Methodology	4
4. Repayment Capacity and Financial Resilience	5
5. Credit Behaviour, Evidence and Source Authority	7
6. Transaction Structure, Collateral and Exposure	9
7. Risk Classification, Safeguards and Uncertainty	11
8. The Role and Boundaries of Artificial Intelligence	13
9. Methodology Development, External Grounding and Testing	14
10. Governance, Limitations and Responsible Institutional Use	16
11. Conclusion	18
Testing Status (as at 20 September 2026)	19
Glossary	20
References and External Sources	23
Revision History	24

1. Executive Overview

The AI Credit Analyst is a purpose-built credit-risk analytical system designed to support institutional consumer and retail lending decisions through a structured and governed methodology.

At its core is an internally developed, proprietary credit-risk framework that defines what information is relevant, how different forms of evidence should be treated, how key credit factors interact, what analytical controls apply, when conclusions may or may not be reached, and where institutional judgement remains necessary.

The framework is organised around the established Five Cs of Credit — Capacity, Character / Credit History, Capital, Collateral and Conditions — but these are treated as interconnected analytical domains rather than simply as headings in an assessment. Each addresses a different dimension of credit risk, while the methodology also considers how those dimensions interact in the context of the proposed transaction.

The AI Credit Analyst incorporates controls governing how source authority, information quality, conflicting evidence, missing information and unresolved uncertainty are treated. Application information, external credit evidence, institutional history, officer observations and controlled calculations are not automatically treated as having the same evidential status.

Its analytical capability does not arise from artificial intelligence alone. Artificial intelligence operates within the platform's broader framework. It performs defined interpretive and analytical functions subject to the methodology, evidence rules and analytical safeguards rather than determining credit policy or lending standards independently.

The AI Credit Analyst is a decision-support system. Its purpose is to identify, analyse and surface material credit risks and other relevant credit considerations rather than to make approval or decline recommendations. Final lending authority remains with the lending institution, which retains responsibility for its own policies, required verification, authorised exceptions and final lending decisions.

The underlying credit-risk principles are not unique to a single jurisdiction. However, lending policies, product structures, regulatory requirements and institutional risk appetites can differ materially. The AI Credit Analyst should therefore not be interpreted as automatically configured or suitable for every lending product or jurisdiction without appropriate review and, where necessary, adaptation and testing.

This paper provides sufficient visibility into the AI Credit Analyst's credit-risk framework to explain its methodological depth, analytical controls, governance principles, and the defined role and limitations of artificial intelligence. Detailed calibration, proprietary decision rules, implementation instructions and other commercially sensitive mechanics remain outside its scope.

2. How the AI Credit Analyst Is Designed

The AI Credit Analyst is designed around the interaction of four principal components: a defined credit risk framework, controlled calculations and data treatment, bounded analytical judgement, and retained institutional authority. Together, these components provide the structure within which individual credit assessments are performed.

2.1 The credit-risk framework

At the core of the platform is an internally developed, proprietary credit risk framework. The framework is built on established, industry-standard and widely recognized credit-risk principles, principally the Five Cs of Credit, and translates those principles into detailed instruction sets, analytical rules and operational controls.

The framework establishes the credit questions that must be addressed and governs how they are addressed. It defines, among other matters:

- the evidence relevant to each area of risk;
- how different information sources may be relied upon;
- required calculations and analytical measures;
- treatment of missing, inconsistent or unreliable information;
- relationships between different credit factors;
- circumstances requiring safeguards, limitations or further review; and
- boundaries between analytical judgement and institutional decision authority.

In practical terms, the credit risk framework, when set out in document form, represents a 212-page compilation of analytical and operational instructions that collectively define how established credit-risk principles are translated into a structured, consistent and governable assessment process.

This framework is an important distinction between the AI Credit Analyst and the use of a general-purpose large language model on its own. A general-purpose model may possess broad knowledge and reasoning capability relevant to lending, but it is not inherently bound to apply a specific credit-risk methodology, evidence hierarchy, calculations, thresholds, rating rules, controls and validation requirements consistently from one assessment to another.

Core Principle

The credit-risk methodology governs the assessment; artificial intelligence operates within that methodology.

Without such a governing framework, the model may still generate plausible credit analysis, but there is no defined basis for ensuring that the same methodology, evidence rules and analytical controls are applied consistently, traceably and within approved boundaries. The distinctive strength of the platform therefore lies in the proprietary credit risk framework within which the AI operates.

2.2 Controlled Calculations and Analytical Inputs

Not every element of credit assessment requires interpretive judgement. Where appropriate, financial calculations are produced through controlled processing and supplied to the analytical framework as authoritative values. The methodology also determines when additional calculations may be made and when a missing value must remain unresolved rather than being estimated without an approved basis.

This preserves an important distinction between calculation and analysis. A calculation establishes a value. Credit analysis determines what that value means when considered alongside the relevant evidence and the proposed transaction.

2.3 Analytical Judgement Within Defined Boundaries

Credit analysis cannot be reduced entirely to arithmetic. Borrower circumstances, evidence quality, repayment history, transaction structure and relationships between risk factors can require contextual interpretation.

The Analyst permits analytical judgement where it is genuinely required, while constraining the information that may support a conclusion and the consequences that particular findings may have. Analytical judgement therefore operates within the methodology rather than independently of it.

2.4 Institutional Authority Remains With the Lender

The AI Credit Analyst is a decision-support system rather than an autonomous lending authority. The lending institution remains responsible for applying its own policies, completing required verification, considering authorised exceptions, satisfying legal and security requirements, and making the final lending decision.

3. The Credit-Risk Methodology

The Five Cs provide the principal structure of the AI Credit Analyst methodology.

3.1 Capacity

Capacity assesses whether the borrower can reasonably meet the proposed repayment obligation while continuing to meet existing financial commitments and household needs. The assessment considers both immediate affordability and the sustainability of the income relied upon for repayment.

3.2 Character / Credit History

Character / Credit History assesses demonstrated repayment reliability and relevant financial conduct using observable evidence rather than assumptions about personality or intent. Relevant evidence may include external credit information, repayment behaviour, delinquency, defaults, disclosure consistency, material legal events and the lender's own repayment experience with the borrower.

3.3 Capital

Capital examines the borrower's broader financial position and resilience, including assets, liabilities, net worth, liquidity and other available financial resources. Its purpose is to understand the strength and accessibility of the borrower's financial position rather than treating balance-sheet wealth alone as sufficient evidence of resilience.

3.4 Collateral

Collateral assesses the quality, usability and sufficiency of security supporting the proposed facility. The methodology distinguishes nominal asset value from the practical protection available to the lender by considering matters such as valuation reliability, enforceability, institutional control, marketability and likely recovery usefulness.

3.5 Conditions

Conditions considers the proposed facility and the circumstances surrounding it, including its purpose, structure, repayment arrangements and other transaction-specific factors. The methodology therefore examines whether the proposed transaction itself is coherent and supportable rather than assessing the borrower independently of the facility being requested.

3.6 Relationships between the Five Cs

The Five Cs are assessed separately so that distinct areas of risk remain visible. They are then considered together so that important relationships between those areas are not lost.

This allows the methodology to distinguish between evidence that genuinely mitigates a particular weakness and favourable evidence that addresses a different credit-risk question.

4. Repayment Capacity and Financial Resilience

Repayment capacity is assessed through multiple related measures rather than a single affordability ratio.

4.1 Debt burden and remaining financial capacity

Debt-service measures provide a useful view of the proportion of income committed to recurring obligations. They do not, however, show how much financial capacity remains after those obligations have been met.

The Analyst therefore considers debt-service burden together with post-facility discretionary income and an internal household-adequacy benchmark. [4] This allows the methodology to recognise circumstances in which a percentage-based affordability measure appears acceptable while the amount remaining to support household needs indicates materially greater pressure.

The household benchmark is an internal analytical control. It is not presented as a measurement of actual household expenditure or as a universal statutory living-cost standard. Its calibration and interaction with other capacity measures remain proprietary.

4.2 Income sustainability

Capacity also considers whether the income relied upon for repayment appears sufficiently dependable. Relevant considerations can include the nature of the income source, stability, employment history, tenure, continuity and other reliable evidence relevant to future repayment.

No single employment characteristic is intended automatically to determine the conclusion. Context matters. For example, relatively short current tenure may have different significance where the wider employment record demonstrates continuity than where the available evidence indicates repeated or unresolved instability.

4.3 The expected post-transaction position

The affordability assessment considers the financial position expected after the proposed transaction. This is particularly important where an application involves refinancing, consolidation or replacement of existing obligations.

The methodology considers the expected effect of the transaction on repayment obligations and cash flow rather than assuming that a refinancing or consolidation label automatically means that affordability improves.

4.4 Financial position and liquidity resilience

Capital complements Capacity by examining broader financial resilience. The Analyst distinguishes overall financial position from the immediate accessibility of financial resources.

A borrower can possess substantial assets while having comparatively limited liquid resources available to absorb a temporary disruption. Conversely, a borrower with a more modest balance sheet may maintain meaningful readily accessible reserves. These are related but distinct credit considerations.

4.5 Liquidity measures are analytical proxies

Where a liquidity-resilience measure is used, it is interpreted according to what the underlying data can legitimately establish. Such a measure may use an income-based proxy rather than measured household expenditure. It should therefore not be interpreted as literal months of actual household living expenses covered.

The measure provides an indicator of accessible financial resilience. It is not a substitute for the Capacity assessment and is not intended to represent actual household expenditure.

Methodology Principle

A metric should answer the credit question for which it was designed and should not be given analytical meaning it does not possess.

5. Credit Behaviour, Evidence and Source Authority

Credit-risk assessment depends not only on what information is available, but also on whether that information is sufficiently current, reliable and relevant to the conclusion being considered.

The Analyst therefore incorporates principles governing source authority, evidence currency, reconciliation and limitations in the available evidence.

5.1 Behavioural credit evidence

Character / Credit History evaluates demonstrated financial behaviour. Relevant evidence can include payment performance, arrears, defaults, collections, credit utilisation, depth of credit experience, recent credit activity, material legal events, disclosure consistency and institutional repayment experience.

The significance of adverse evidence is considered in context. Recency, severity, frequency, persistence, resolution and subsequent conduct can all affect its relevance.

5.2 Evidence does not all perform the same function

The Analyst distinguishes between evidence that may anchor an assessment and evidence that may provide support, qualification or adverse context.

Where an eligible current external credit score provides the applicable Character / Credit History anchor, other evidence can still materially inform the assessment without automatically replacing the role assigned to that anchor. Where no eligible anchor exists, the methodology can apply a different evidence pathway using the information legitimately available.

This distinction allows multiple sources to contribute to the assessment without treating them as interchangeable. Detailed score bands, fallback rules and rating-conversion mechanics remain proprietary.

5.3 External and internal repayment information

External credit information and institutional repayment history can both be relevant while serving different purposes. External information can provide a wider view of obligations and repayment behaviour across lenders. Internal institutional history can provide direct evidence of the institution's own experience with the borrower.

The methodology is designed to use both appropriately while avoiding duplicate analytical weight for the same underlying behaviour.

5.4 Information currency

Credit information can lose relevance with time. The methodology therefore considers whether external information remains sufficiently current for the purpose for which it is being used. Older information may retain historical relevance, but it is not automatically treated as equivalent to current evidence.

Detailed currency periods and their internal consequences are not reproduced in this paper.

5.5 Reconciliation of conflicting evidence

Different sources do not always agree. The Analyst includes controls for identifying material differences between application information, institutional records and permitted external sources.

For example, where an external source identifies a material liability that is not reflected in the information used for affordability analysis, the difference may require reconciliation before the resulting affordability position can be relied upon.

The framework considers the reliability and relevance of the sources, the significance of the discrepancy, and whether the matter can be addressed analytically or requires institutional clarification. It does not assume that every discrepancy can be resolved automatically.

5.6 Legal and adverse events are considered in context

The existence of a material adverse event does not, by itself, necessarily establish its current credit significance. Relevant considerations can include recency, severity, status, resolution, subsequent conduct and disclosure consistency.

This allows historical evidence to retain appropriate relevance without automatically treating every past event as an unresolved present condition.

5.7 Absence of evidence is not favourable or adverse evidence

The Analyst distinguishes among:

- adverse behaviour that has been established;
- favourable behaviour that has been established;
- information that is unavailable; and
- information that is too limited to support a reliable conclusion.

The absence of adverse information is not automatically treated as proof of satisfactory behaviour. Equally, unavailable information is not automatically converted into an adverse conclusion.

6. Transaction Structure, Collateral and Exposure

The AI Credit Analyst assesses the borrower and the proposed facility together. The transaction itself can create, reduce or reshape risk depending on its purpose, structure, term, repayment arrangements and security.

6.1 Transaction coherence

The methodology assesses whether the essential elements of the proposed facility reasonably fit together.

- clarity of the proposed use of funds;
- whether the selected facility type fits that purpose;
- whether the amount is broadly proportionate;
- whether the repayment term is reasonable for what is being financed; and
- whether refinancing, consolidation, settlement or additional-funds arrangements are sufficiently clear.

Not every minor information gap is treated as adverse. Clarification becomes important where the missing or inconsistent information is material to understanding the transaction.

6.2 Policy and structural considerations

Some facility characteristics can create policy or structural issues that are distinct from ordinary affordability. For example, borrower age at scheduled loan maturity may be relevant where institutional policy, product requirements or insurance eligibility depend upon it.

This is separate from collateral-specific measures such as the age of a financed asset at scheduled maturity. Such considerations are not presented as universal rules. Their relevance depends on the product, institution and applicable jurisdictional or insurer requirements.

6.3 Collateral quality and usability

Collateral is not assessed simply by comparing a stated asset value with the amount being financed. The methodology can consider valuation reliability, enforceability, lender control, charge position, ownership, marketability, depreciation risk and likely ease of realisation.

A high-value asset does not necessarily provide strong lender protection if the security cannot be effectively established or the asset would be difficult to realise. Collateral quality therefore matters alongside collateral quantity.

6.4 Primary repayment and secondary recovery

Collateral performs a different function from repayment capacity. The primary question is whether the borrower can reasonably meet the proposed repayments. Collateral addresses the lender's recovery position if repayment fails.

The Analyst does not permit strong security automatically to substitute for fundamental primary-repayment support. [1][2][3] The exact safeguards governing severe primary-repayment weakness remain proprietary.

6.5 Facility-level exposure and institution-level exposure

The Analyst distinguishes among related but different exposure concepts.

- Total Exposure - a proposed-facility measure.
- Facility Unsecured Exposure - the unsecured portion of the proposed facility after eligible collateral treatment.
- Institution Unsecured Exposure Limit - the AI Credit Analyst's calculated recommended limit for the member or applicant's total unsecured exposure with the institution. This corresponds to the methodology's member-level Maximum Recommended Exposure concept.

These measures should not be confused. A proposed facility can appear acceptable when compared with the calculated Institution Unsecured Exposure Limit while the institution still needs to determine the member or applicant's actual aggregate unsecured position after taking existing exposure into account.

The calculated Institution Unsecured Exposure Limit is also distinct from any formally adopted institutional policy limit.

7. Risk Classification, Safeguards and Uncertainty

Risk classification within the AI Credit Analyst is constrained by the methodology rather than produced as unrestricted AI judgement.

7.1 Ratings follow the analytical findings

Section classifications and the overall assessment are linked to the evidence and analytical conclusions established within the framework. Material weaknesses, limitations in the evidence and defined safeguards are required to retain their methodological consequences rather than being overridden simply because the case appears attractive overall.

7.2 Compensating factors have limits

Credit analysis often involves strengths that mitigate weaknesses. The Analyst recognises relevant compensating factors, but it does not treat all favourable evidence as interchangeable. A favourable factor can compensate only where it addresses the underlying risk question in a manner permitted by the methodology.

This prevents a simple approach in which enough positive factors can automatically cancel any type of weakness.

7.3 Override safeguards

Certain primary-repayment conditions can have consequences beyond ordinary weighted assessment. Where the methodology identifies sufficiently serious primary-repayment weakness, a safeguard can prevent the assessment from continuing through the ordinary risk-classification path.

The purpose is to ensure that fundamental primary-repayment risk is not diluted solely by stronger secondary factors. The exact triggers, weights and calibration remain proprietary.

7.4 Not Available has more than one cause

A material analytical conclusion can be Not Available for different reasons. One pathway arises where reliable evidence is insufficient to support the conclusion. A separate pathway arises where a methodological safeguard prevents completion of the ordinary assessment required to produce that rating.

In the second case, the underlying facts may still be known and reportable. For example, collateral information may remain available even though the methodology does not permit completion of an ordinary Collateral rating after a primary-repayment override.

Where a mandatory section cannot be completed, the overall composite rating is Not Available. Not Available is therefore not the same as a poor rating, and it does not by itself constitute an institutional decline decision.

7.5 Material uncertainty remains part of the assessment

The framework requires unresolved material matters to remain visible through an appropriate mechanism. Depending on the circumstances, this may include:

- a limitation;
- a verification requirement;
- a required action or condition;
- an unavailable conclusion; or
- another defined control outcome.

The objective is to avoid converting uncertainty into false precision.

7.6 Classification is decision support, not a default probability

The Analyst's classifications communicate the methodology's assessment of the credit-risk position under the evidence and rules applied. They are not, by themselves, approval decisions and are not represented as statistically calibrated probabilities of future default unless separate empirical evidence establishes such a relationship.

8. The Role and Boundaries of Artificial Intelligence

Artificial intelligence is an important component of the AI Credit Analyst, but it is not the credit-risk methodology.

8.1 Where AI adds value

AI is used primarily where credit assessment requires interpretation and synthesis rather than calculation alone. Within the governing framework, it can assess the significance of relevant evidence, identify relationships between credit factors, recognise inconsistencies, interpret material strengths and weaknesses, and integrate the findings into a coherent credit-risk assessment.

These functions are valuable because many underwriting questions cannot be resolved through arithmetic alone.

8.2 The methodology constrains the AI

AI is not intended to create its own lending methodology, invent unsupported facts, replace authoritative calculated values, disregard defined controls, conceal uncertainty through unsupported assumptions, independently change established policy treatment or make the final lending decision.

The methodology establishes the permissible analytical boundaries. Different implementation mechanisms may then be used to execute or enforce those requirements.

8.3 Defined controls do not eliminate execution risk

Prescribed analytical boundaries and controls are intended to constrain AI behaviour, but they do not guarantee perfect execution. Incorrect interpretation, omission or inconsistent application can still occur.

Structured testing, traceability, change control and authorised human review therefore remain necessary components of responsible use. [5][6][7]

8.4 AI does not replace source authority or institutional judgement

The availability of an AI inference does not make unsupported information reliable. Evidence must still satisfy the methodological requirements governing source authority and relevance. Likewise, final lending authority remains with the institution.

The analytical capability of the Analyst therefore comes from the interaction of methodology, controlled calculations, evidence rules, safeguards, AI interpretation and human institutional oversight - not from the AI component in isolation.

9. Methodology Development, External Grounding and Testing

The AI Credit Analyst has been developed through an iterative methodology-design process centred on practical credit-risk problems.

9.1 Development begins with the credit problem

Methodology development begins by identifying the underlying underwriting question. Examples include determining:

- what evidence is required to support a particular conclusion;
- how different information sources should interact;
- when a transaction creates risk not captured by general borrower metrics;
- when historical adverse evidence remains relevant;
- what characteristics make collateral genuinely useful to the lender;
- when missing information prevents a reliable conclusion; and
- when favourable evidence legitimately compensates for a weakness.

The starting point is therefore the credit-risk problem, not the technical capabilities of the AI.

9.2 Review of the methodology

Where appropriate, proposed analytical treatments are considered against established supervisory guidance, recognised lending practice and relevant credit-risk principles. [1][2][3][4] This provides a basis for assessing whether the assumptions, relationships and intended consequences of a proposed treatment are reasonable.

External grounding does not determine the AI Credit Analyst's exact implementation. The methodology used to operationalise those principles - including internal calibration, weighting, interaction rules, safeguards and exception treatment - remains proprietary.

References to external standards support the relevant underlying principle; they do not imply approval, certification, validation or endorsement of the AI Credit Analyst by the referenced organisation.

9.3 Implementation testing

A reasonable methodology does not itself establish that the implemented system applies it correctly. CreditVantage AI therefore uses structured testing to examine implemented behaviour.

Testing can include:

- representative lending scenarios;
- boundary conditions;
- missing and contradictory information;
- combinations of favourable and adverse evidence;
- safeguard and unavailable-outcome cases;
- repeated-input testing;
- regression testing following material changes; and
- consistency checks across related implementation paths.

The purpose is to assess whether the system produces the intended methodological consequences for the cases and configurations examined.

9.4 Testing conclusions are bounded

Different forms of testing establish different things. A calculation test can establish whether a defined mathematical rule produces the expected result. A scenario test can establish whether the implemented system applies the intended methodology correctly to that scenario.

Regression testing can determine whether previously tested behaviour remains intact following change. Repeated-input testing can provide evidence about execution consistency under the conditions examined.

These forms of testing do not, by themselves, establish how accurately the methodology predicts subsequent default, delinquency or credit loss.

Accordingly, CreditVantage AI distinguishes among:

- Methodology review - whether the analytical treatment itself is reasonable and appropriately supported.
- Implementation testing - whether the implemented system applies the approved treatment as intended.
- Outcome validation - whether assessments demonstrate an empirically established relationship with subsequent real-world credit performance.

Because the term validated can imply different types and levels of evidence, any statement regarding validation should be limited to the specific evidence actually available.

9.5 Change control forms part of methodology governance

Credit methodology, technology and institutional requirements can evolve. Material changes are therefore treated as governance events requiring evaluation rather than being assumed automatically to improve the system.

Relevant functionality should be retested where appropriate, with regression testing used to identify unintended effects on previously established behaviour. [5][7] A separate dated Testing Status disclosure identifies the testing evidence applicable to the relevant AI Credit Analyst configuration and the limitations of that evidence.

10. Governance, Limitations and Responsible Institutional Use

The AI Credit Analyst is designed as institutional decision support operating within defined methodological, governance and human-oversight boundaries.

10.1 Institutional responsibilities

The lending institution retains responsibility for the final lending decision. It also remains responsible for applying its own policies, completing required verification, addressing authorised exceptions, satisfying legal or security requirements and considering relevant information outside the AI Credit Analyst assessment where necessary.

10.2 CreditVantage AI governance responsibilities

CreditVantage AI retains responsibility for maintaining and governing its methodology and implemented analytical framework. This includes, as applicable:

- maintaining approved methodology documentation;
- controlling substantive methodology versions;
- reviewing material methodological or technology changes;
- investigating identified methodology or system issues;
- implementing corrective changes where necessary;
- performing relevant regression testing following material change; and
- maintaining records of substantive methodology changes and testing activity.

These responsibilities reflect the current governance model and should not be interpreted as implying independent external validation or governance structures where none have been established.

10.3 Traceability of the analytical basis

A credit conclusion is more useful when its basis can be reviewed. The Analyst maintains structured analytical information supporting the assessment, including relevant calculations, material findings, limitations, source-treatment decisions, safeguards and rating outcomes.

The objective is to provide an inspectable record of the evidence and analytical conclusions used by the framework - not access to an AI model's private internal reasoning.

10.4 Scope and jurisdictional adaptability

The underlying methodology is substantially based on credit-risk principles that can apply across lending markets. However, product structures, institution-specific policies, regulatory obligations, risk appetites and local market conditions can differ.

The framework therefore distinguishes between:

- broadly applicable credit-risk principles;
- CreditVantage AI's proprietary analytical methodology; and
- institution- or jurisdiction-specific requirements.

Adaptation to a different institution, product or jurisdiction may require configuration, methodological review, development or additional testing. The existence of a broadly applicable principle should not be interpreted as evidence that every implementation is already suitable for every market.

10.5 Data-use boundaries

The Analyst is designed to use financial, credit and risk information relevant to the assessment while limiting the analytical use of identifying information that is not required for the credit-risk purpose. This is a principle of analytical purpose limitation.

It does not imply that identifying information is never processed, transmitted or retained elsewhere in the broader operational environment. Detailed privacy, security and data-handling arrangements are appropriately addressed through the applicable privacy, security and institutional due-diligence documentation.

10.6 Defined scope limitations

The Analyst does not replace:

- identity verification;
- customer due diligence;
- legal review;
- independent valuation where required;
- security perfection;
- institution-specific policy approval;
- regulatory functions outside the scope of the credit assessment; or
- final lending governance.

10.7 Responsible institutional reliance

The Analyst is intended to provide a more structured and analytically disciplined basis for credit assessment while preserving institutional accountability.

Its outputs should be considered together with the underlying application information, outstanding verification requirements, applicable policy, legal and regulatory obligations, and authorised human judgement.

Institutional Accountability

The AI Credit Analyst provides a structured analytical process; the lending institution retains responsibility for the lending decision.

11. Conclusion

The AI Credit Analyst has been developed as a purpose-built institutional consumer and retail credit-risk analytical system in which artificial intelligence operates within a defined methodology rather than serving as the methodology itself.

Its capability arises from the interaction of structured credit principles, controlled calculations, defined evidence treatment, transaction analysis, relationships between the Five Cs, analytical safeguards, treatment of uncertainty and retained institutional oversight.

Artificial intelligence contributes interpretation and synthesis where those functions add value, while the methodology establishes the questions to be answered, the evidence that may support them, the applicable analytical boundaries and the consequences of material findings.

The methodology has been developed iteratively through practical underwriting problem-solving, external consideration of established credit-risk principles, proprietary methodological design and structured testing of implemented behaviour.

The Analyst is intended to provide lenders with a transparent and disciplined basis for evaluating consumer and retail credit risk while preserving the professional judgement, policy authority and institutional accountability required for responsible lending.

Testing Status (as at 20 September 2026)

Status date 20 September 2026

Configuration covered	AI Credit Analyst production methodology and implementation baseline current as at 17 September 2026
------------------------------	--

CreditVantage AI has undertaken internally conducted testing to assess whether the implemented system applies the approved credit-risk methodology as intended.

Testing completed to date has included calculation and rule verification, controlled credit scenarios, boundary conditions, missing and contradictory information, safeguard and Not Available outcomes, repeated-input consistency testing, regression testing following material changes, and consistency checks across related implementation paths.

Testing is an ongoing governance activity. Additional testing is undertaken where material changes to methodology, analytical controls, implementation logic or underlying technology require reassessment.

The testing completed to date has been conducted internally by CreditVantage AI and should not be interpreted as independent external validation.

The current testing programme assesses methodological and implementation behaviour. It does not establish an empirically validated relationship between AI Credit Analyst assessments and subsequent borrower default, delinquency or credit-loss outcomes. No such outcome-validation claim is made on the basis of the testing completed to date.

Glossary

Term	Definition
Capacity	The assessment of whether the borrower can reasonably meet the proposed repayment obligation while continuing to meet existing financial commitments and household needs, taking account of both affordability and income sustainability.
Character / Credit History	The assessment of demonstrated repayment behaviour and relevant financial conduct using available internal and external credit evidence.
Capital	The assessment of the borrower's broader financial position and resilience, including assets, liabilities, net worth, liquidity and other available financial resources.
Collateral	Assets or other security supporting a credit facility and providing a potential secondary source of recovery if repayment fails. The AI Credit Analyst considers not only value, but also factors such as enforceability, control, marketability and likely recovery usefulness.
Conditions	The assessment of the proposed transaction and relevant circumstances surrounding it, including purpose, structure, repayment arrangements and other transaction-specific considerations.
Debt-Service Ratio (DSR)	A measure of the proportion of income committed to recurring debt obligations. The AI Credit Analyst uses debt-service measures as part of the Capacity assessment rather than as a complete measure of affordability on their own.
Discretionary Income	Income remaining after the debt obligations recognised by the methodology have been taken into account. Within the AI Credit Analyst methodology, it is used alongside other Capacity measures to assess remaining financial capacity.
Household-Adequacy Benchmark	An internal analytical benchmark used to evaluate whether post-facility discretionary income appears sufficient relative to household circumstances. It is not a measurement of actual household expenditure or a statutory living-cost standard.
Liquidity Resilience	The extent to which readily accessible financial resources may provide support against short-term financial disruption. It is distinct from overall net worth or balance-sheet wealth.
Source Authority	The methodological principle governing how much reliance may appropriately be placed on a particular information source for a particular credit conclusion. Different sources may perform different analytical functions.
Anchoring Evidence	Evidence assigned a primary methodological role in establishing a particular assessment or rating. Other evidence may remain supportive, qualifying or adverse without necessarily replacing the anchor.

Term	Definition
Primary Repayment	Repayment expected to arise from the borrower's ordinary income and financial capacity. It is the principal basis on which repayment ability is assessed.
Secondary Recovery	Recovery available to the lender from collateral or other security if primary repayment fails. Strong secondary recovery does not automatically substitute for adequate primary repayment support.
Analytical Safeguard	A methodological control that constrains how an assessment may proceed when specified risk, evidential or structural conditions arise.
Override Review	A controlled assessment state triggered by sufficiently serious primary-repayment concerns under the AI Credit Analyst methodology. Exact trigger conditions and calibration are proprietary.
Not Available	An outcome indicating that a required analytical conclusion cannot validly be completed. This may arise because reliable evidence is insufficient or because a methodological safeguard prevents completion of the ordinary assessment path. Not Available is not equivalent to a poor rating or an automatic decline decision.
Total Exposure	A proposed-facility measure representing the total exposure associated with the facility being assessed. It should not be confused with the member/applicant's aggregate exposure to the institution.
Facility Unsecured Exposure	The portion of the proposed facility that remains unsecured after the methodology's eligible collateral treatment.
Institution Unsecured Exposure Limit	The AI Credit Analyst's calculated recommended limit for the member/applicant's total unsecured exposure with the institution. It represents the report-facing expression of the member-level Maximum Recommended Exposure methodology and is distinct from both Facility Unsecured Exposure and any policy limit formally adopted by the institution.
Maximum Recommended Exposure (MRE)	The underlying AI Credit Analyst methodology used to determine the recommended maximum member/applicant-level unsecured institutional exposure. It is not the same as the Total Exposure of the proposed facility.
Methodology Review	Evaluation of whether the assumptions, analytical relationships and intended consequences of a credit-risk treatment are reasonable and appropriately supported.
Implementation Testing	Testing performed to determine whether the implemented system applies the approved methodology as intended for the scenarios and configurations examined.
Outcome Validation	Empirical evaluation of whether credit assessments demonstrate an established relationship with subsequent real-world outcomes such as delinquency, default or credit loss.
Calibration	The setting of thresholds, ranges, weights or other parameters used to operationalise a methodology. Detailed AI Credit Analyst calibration is proprietary.

Term	Definition
Decision Support	Analytical assistance provided to authorised institutional decision-makers. The AI Credit Analyst supports the lending process but does not itself exercise final lending authority.

References and External Sources

The external references used in this paper support selected underlying credit-risk, lending and AI-governance principles. They do not constitute validation, certification or endorsement of CreditVantage AI, nor do they determine CreditVantage AI's proprietary methodology, calibration or implementation.

1. Basel Committee on Banking Supervision. Principles for the Management of Credit Risk. Bank for International Settlements, 30 April 2025. Primary international reference for sound credit-granting processes, repayment capacity, transaction-level assessment, collateral and credit-risk controls. Official source (accessed 20 September 2026).
2. European Banking Authority. Guidelines on Loan Origination and Monitoring (EBA/GL/2020/06). 2020. Supporting European supervisory guidance on creditworthiness assessment, prudent loan origination, governance and monitoring. Official source (accessed 20 September 2026).
3. Office of the Comptroller of the Currency. Comptroller's Handbook: Retail Lending. October 2021. United States retail-lending supervisory guidance used selectively for retail credit-risk principles. Official source (accessed 20 September 2026).
4. Consumer Financial Protection Bureau. Regulation Z, 12 CFR 1026.43 - Minimum Standards for Transactions Secured by a Dwelling, including Official Interpretations. United States mortgage-specific authority cited narrowly for recognised use of debt-to-income and residual-income measures in repayment-capacity assessment. Official source (accessed 20 September 2026).
5. National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1. 26 January 2023. Voluntary, non-sector-specific framework for AI risk governance across design, deployment, use and evaluation. AI RMF 1.0 remains under revision as of September 2026. Official source (accessed 20 September 2026).
6. National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, NIST AI 600-1. 26 July 2024. Companion resource addressing generative-AI risks and risk-management considerations. Official source (accessed 20 September 2026).
7. National Institute of Standards and Technology. AI Risk Management Framework resources and AI Resource Center. Operational resources supporting AI risk management, testing, evaluation, verification and validation practices. Official source (accessed 20 September 2026).

Reference Scope

External sources cited in this paper are used to demonstrate the established basis of selected credit-risk, lending and AI-governance principles. Jurisdiction-specific regulatory sources are cited only for the particular principle illustrated and should not be interpreted as establishing regulatory applicability to the AI Credit Analyst in another jurisdiction. CreditVantage AI's proprietary methodology, calibration and implementation remain independently developed and are not represented as having been approved, validated or endorsed by the referenced organisations.

Revision History

Version	Date	Status	Summary of Changes
1.0	20 September 2026	Initial Public Release	Initial public release, including the main methodology and governance paper, Testing Status, glossary, verified references and document controls.