Methodology summary

Each vendor is scored on 12 axes drawn from US AI governance frameworks: HIPAA BAA / DPA availability, training-data opt-out default, US data residency option, SOC 2 Type II report, ISO/IEC 42001 attestation, NIST AI RMF self-attestation, Colorado AI Act SB 24-205 readiness, HHS-OCR Section 1557 readiness, FRB SR 11-7 readiness, ABA Formal Opinion 512 readiness, public subprocessor list, and a 1–5 trust-center maturity sub-score.

Status values are **yes** (control in place and documented), **partial** (documented but conditional or tier-gated), **no** (control absent or undocumented), or **na**(axis does not apply to the vendor's deployment scope — e.g. Section 1557 for a general-purpose foundation model). N/A axes are excluded from the composite denominator so vendors are not penalized for axes outside their scope. The composite is sector-weighted: healthcare-vertical vendors carry 2× weight on Section 1557, legal-vertical on ABA Op 512, banking-vertical on SR 11-7, with a 1.5× BAA multiplier for any regulated sector. Trust-center maturity contributes 10% of the final composite on top of the 12-axis average.

Every cell is source-cited to public vendor documentation (trust portals, BAAs/DPAs, SOC report cover pages, model cards, and published methodology). This is an EFROS research artifact, not legal or compliance advice. Posture changes frequently — verify with the vendor's trust center before contract.

Full ranking — all 30 vendors

Sorted by sector-weighted composite (descending). Per-axis cells show yes / partial / no / na status.

#	Vendor	Category	Composite	Grade	BAA	Opt-out	US Res	SOC 2	ISO 42001	NIST AI	CO AI	§1557	SR 11-7	ABA 512	Subproc	TC
1	**Abridge**	Healthcare	87	**A**	Yes	Yes	Yes	Yes	Partial	Partial	Partial	Yes	N/A	N/A	Yes	5/5
2	**Thomson Reuters CoCounsel**	Legal	80	**B**	Yes	Yes	Yes	Yes	No	Partial	Partial	N/A	N/A	Yes	Yes	4/5
3	**FICO Falcon Fraud Manager + FICO Score AI**	Banking	80	**B**	Yes	Yes	Yes	Yes	No	Partial	Partial	N/A	Yes	N/A	Yes	4/5
4	**Lexis+ AI**	Legal	76	**B**	Yes	Yes	Yes	Yes	No	Partial	No	N/A	N/A	Yes	Yes	4/5
5	**Westlaw Precision AI**	Legal	76	**B**	Yes	Yes	Yes	Yes	No	Partial	No	N/A	N/A	Yes	Yes	4/5
6	**Microsoft 365 Copilot**	Productivity	75	**B**	Yes	Yes	Yes	Yes	Partial	Partial	Partial	Partial	Partial	Partial	Yes	5/5
7	**Harvey**	Legal	74	**B**	Yes	Yes	Yes	Yes	No	Partial	Partial	N/A	N/A	Yes	Partial	3/5
8	**Zest AI**	Banking	74	**B**	Yes	Yes	Yes	Yes	No	Partial	Partial	N/A	Yes	N/A	Partial	3/5
9	**Upstart**	Banking	74	**B**	Yes	Yes	Yes	Yes	No	Partial	Partial	N/A	Yes	N/A	Partial	3/5
10	**Suki AI**	Healthcare	72	**B**	Yes	Yes	Yes	Yes	No	Partial	Partial	Partial	N/A	N/A	Yes	4/5
11	**Nuance DAX Copilot (Microsoft)**	Healthcare	70	**B**	Yes	Yes	Yes	Yes	No	Partial	No	Partial	N/A	N/A	Yes	5/5
12	**Salesforce Einstein / Agentforce**	Productivity	69	**C**	Yes	Yes	Yes	Yes	No	Partial	No	Partial	Partial	N/A	Yes	5/5
13	**Glean**	Productivity	69	**C**	Yes	Yes	Yes	Yes	No	Partial	No	N/A	N/A	N/A	Yes	4/5
14	**Arctic Wolf**	security-mssp	69	**C**	Yes	Yes	Yes	Yes	No	Partial	No	N/A	N/A	N/A	Yes	4/5
15	**Huntress**	security-mssp	69	**C**	Yes	Yes	Yes	Yes	No	Partial	No	N/A	N/A	N/A	Yes	4/5

#	Vendor	Category	Composite	Grade	BAA	Opt-out	US Res	SOC 2	ISO 42001	NIST AI	CO AI	§1557	SR 11-7	ABA 512	Subproc	TC
16	**eSentire**	security-mssp	69	**C**	Yes	Yes	Yes	Yes	No	Partial	No	N/A	N/A	N/A	Yes	4/5
17	**Sophos**	security-mssp	69	**C**	Yes	Yes	Yes	Yes	No	Partial	No	N/A	N/A	N/A	Yes	4/5
18	**Unit21**	Banking	68	**C**	Yes	Yes	Yes	Yes	No	Partial	No	N/A	Partial	N/A	Yes	4/5
19	**Ironclad AI**	Legal	63	**C**	Yes	Yes	Yes	Yes	No	No	No	N/A	N/A	Partial	Yes	4/5
20	**Anthropic Claude**	Foundation	58	**C**	Partial	Yes	Partial	Yes	No	Partial	No	N/A	N/A	N/A	Yes	4/5
21	**Google Gemini for Workspace**	Foundation	58	**C**	Partial	Partial	Yes	Yes	No	Partial	No	N/A	N/A	N/A	Yes	4/5
22	**Hummingbird**	Banking	56	**C**	Yes	Yes	Yes	Yes	No	No	No	N/A	Partial	N/A	Partial	3/5
23	**OpenAI ChatGPT & API**	Foundation	53	**D**	Partial	Partial	Partial	Yes	No	Partial	No	N/A	N/A	N/A	Yes	4/5
24	**ConnectWise**	security-mssp	50	**D**	Partial	Yes	Partial	Yes	No	No	No	N/A	N/A	N/A	Yes	3/5
25	**Spellbook**	Legal	45	**D**	Yes	Yes	Partial	Partial	No	No	No	N/A	N/A	Partial	Partial	2/5
26	**Heidi Health**	Healthcare	45	**D**	Yes	Yes	Partial	Partial	No	No	No	Partial	N/A	N/A	Partial	2/5
27	**Notion AI**	Productivity	33	**F**	No	Partial	No	Yes	No	No	No	N/A	N/A	N/A	Yes	3/5
28	**Meta Llama**	Foundation	25	**F**	No	Yes	Yes	No	No	No	No	N/A	N/A	N/A	No	2/5
29	**Otter.ai**	Productivity	25	**F**	No	Partial	No	Yes	No	No	No	N/A	N/A	N/A	Partial	2/5
30	**Perplexity AI**	Foundation	19	**F**	No	Partial	No	Partial	No	No	No	N/A	N/A	N/A	Partial	2/5

BAA = HIPAA Business Associate Agreement; Opt-out = training-data opt-out default; US Res = US data residency option; SOC 2 = SOC 2 Type II; ISO 42001 = ISO/IEC 42001 attestation; NIST AI = NIST AI RMF self-attestation; CO AI = Colorado AI Act readiness; §1557 = HHS-OCR Section 1557 readiness; SR 11-7 = FRB SR 11-7 readiness; ABA 512 = ABA Formal Opinion 512 readiness; Subproc = subprocessor list public; TC = trust-center maturity (1–5).

12-AXIS SCORING

Axis	Status	Note	Source
BAA / DPA available	Yes	Abridge signs BAAs for all enterprise customers.	Abridge Trust
Training-data opt-out	Yes	Customer audio and notes not used for general model training. Tenant isolation enforced.	Abridge Trust
US data residency option	Yes	Abridge hosted on US infrastructure. US data residency standard for US customers.	Abridge Trust
SOC 2 Type II report	Yes	Abridge holds SOC 2 Type II.	Abridge Trust
ISO/IEC 42001 attestation	Partial	Abridge has publicly indicated ISO/IEC 42001 alignment work in progress. Certification not yet posted as of May 2026.	Abridge governance documentation
NIST AI RMF self-attestation	Partial	Abridge publishes a Responsible AI framework mapped against NIST AI RMF functions.	Abridge Responsible AI
Colorado AI Act readiness	Partial	Abridge has publicly engaged on the Colorado AI Act deployer-responsibility model; product documentation addresses high-risk classification.	Abridge customer documentation
HHS-OCR Section 1557 readiness	Yes	Abridge has publicly addressed Section 1557 algorithmic non-discrimination — bias testing, model card publication, ongoing monitoring documentation.	Abridge Section 1557 documentation
FRB SR 11-7 readiness	N/A	Healthcare-vertical positioning.	Abridge positioning
ABA Formal Op 512 readiness	N/A	Healthcare-vertical positioning.	Abridge positioning
Subprocessor list public	Yes	Abridge subprocessor list public via trust center.	Abridge Trust
Trust-center maturity	5/5	Abridge's trust center is one of the most mature in clinical AI — public Responsible AI framework, Section 1557 documentation, model cards, subprocessor transparency.	Abridge Trust

DEEP DIVE

Abridge is one of the very few clinical AI vendors that has directly engaged the Section 1557 algorithmic non-discrimination requirement — most vendors in the category punt this to deployer responsibility. Combined with strong platform fundamentals (BAA, residency, SOC 2) and a mature trust center, Abridge has the cleanest US healthcare AI governance posture in the index.

STRENGTHS

- Direct Section 1557 algorithmic non-discrimination engagement
- Public Responsible AI framework + model cards
- BAA, US residency, SOC 2 Type II
- Mature trust center

WEAKNESSES

- ISO/IEC 42001 in progress, not yet certified
- Pricing typically higher than Microsoft DAX Copilot at scale

BEST USE CASE

Health systems prioritizing best-in-class clinical AI governance — particularly those with active OCR scrutiny on Section 1557 or those running quality programs that benefit from public model card documentation.

AVOID WHEN

Microsoft 365-standardized health systems where DAX Copilot's M365/Azure inheritance and EHR integration breadth fit existing IT operations better.

Last reviewed: 2026-05-13 · Homepage: https://www.abridge.com · Trust center: https://www.abridge.com/trust

---

12-AXIS SCORING

Axis	Status	Note	Source
BAA / DPA available	Yes	CoCounsel is covered under Thomson Reuters' enterprise data-handling agreements. BAA scope addressed for firms with PHI in matter content.	Thomson Reuters Trust Center
Training-data opt-out	Yes	CoCounsel does not train models on customer data. Tenant isolation enforced.	Thomson Reuters CoCounsel Privacy
US data residency option	Yes	US data residency available for enterprise customers.	Thomson Reuters Trust Center
SOC 2 Type II report	Yes	Thomson Reuters Cloud Platform (which hosts CoCounsel) holds SOC 2 Type II and ISO 27001.	Thomson Reuters Trust Center
ISO/IEC 42001 attestation	No	No ISO/IEC 42001 attestation as of May 2026.	Public posture review
NIST AI RMF self-attestation	Partial	Thomson Reuters publishes AI Principles and governance documentation; no formal NIST AI RMF self-attestation.	Thomson Reuters AI Principles
Colorado AI Act readiness	Partial	Thomson Reuters documents the deployer responsibility model under Colorado AI Act.	Thomson Reuters customer documentation
HHS-OCR Section 1557 readiness	N/A	Legal-vertical positioning.	Thomson Reuters CoCounsel positioning
FRB SR 11-7 readiness	N/A	Legal-vertical positioning.	Thomson Reuters CoCounsel positioning
ABA Formal Op 512 readiness	Yes	Thomson Reuters publishes ABA Op 512 alignment documentation specific to CoCounsel deployment.	Thomson Reuters CoCounsel ABA Op 512 documentation
Subprocessor list public	Yes	Subprocessor list published as part of Thomson Reuters Cloud Platform terms.	Thomson Reuters Subprocessors
Trust-center maturity	4/5	Thomson Reuters Trust Center is mature for cloud-platform compliance; AI-specific governance for CoCounsel is documented but less granular than the platform compliance.	Thomson Reuters Trust Center

DEEP DIVE

CoCounsel benefits from the parent Thomson Reuters compliance stack — well above what most legal-vertical AI vendors offer on their own. Tight integration with Westlaw and Practical Law content reduces hallucination risk on legal research workflows. The governance posture is more mature than Harvey on platform fundamentals; the workflow differentiation depends on firm preference.

STRENGTHS

- Inherits Thomson Reuters Cloud Platform compliance stack
- ABA Op 512 alignment documented
- Tight integration with Westlaw / Practical Law — citation grounding
- Mature subprocessor transparency

WEAKNESSES

- No ISO/IEC 42001
- No formal NIST AI RMF self-attestation
- Pricing structure is more complex than per-seat alternatives

BEST USE CASE

Firms already standardized on Westlaw and Practical Law, where CoCounsel's content integration delivers operational value beyond raw generative drafting.

AVOID WHEN

Firms standardized on Lexis content — CoCounsel's research integration value depends on Westlaw/Practical Law alignment.

Last reviewed: 2026-05-13 · Homepage: https://www.thomsonreuters.com/en/artificial-intelligence/cocounsel.html

12-AXIS SCORING

Axis	Status	Note	Source
BAA / DPA available	Yes	FICO signs DPAs / data-handling agreements for enterprise customers. BAA available where PHI exposure exists in customer datasets.	FICO Trust
Training-data opt-out	Yes	Customer transaction data is processed under contracted purpose limitation; not used for cross-customer model training without explicit consortium opt-in.	FICO Trust
US data residency option	Yes	US data residency available for US bank customers. FICO operates US-region data centers + AWS GovCloud for federal-aligned deployments.	FICO Trust
SOC 2 Type II report	Yes	FICO holds SOC 2 Type II, ISO 27001, FedRAMP. Most banks have FICO compliance documentation already on file.	FICO Trust
ISO/IEC 42001 attestation	No	No ISO/IEC 42001 attestation as of May 2026.	Public posture review
NIST AI RMF self-attestation	Partial	FICO publishes a Responsible AI framework with explicit NIST AI RMF mapping; no formal self-attestation document.	FICO Responsible AI
Colorado AI Act readiness	Partial	FICO has publicly engaged on the Colorado AI Act and deployer-responsibility documentation for credit decisioning customers.	FICO customer documentation
HHS-OCR Section 1557 readiness	N/A	Banking-vertical positioning.	FICO positioning
FRB SR 11-7 readiness	Yes	FICO model documentation is the reference SR 11-7 validation packet in the credit-scoring industry. Validation reports, conceptual soundness reviews, ongoing performance monitoring all packaged for examiner review.	FICO SR 11-7 documentation packet
ABA Formal Op 512 readiness	N/A	Banking-vertical positioning.	FICO positioning
Subprocessor list public	Yes	FICO subprocessor list available to enterprise customers.	FICO Trust
Trust-center maturity	4/5	Mature compliance documentation, broad certificate library, SR 11-7-grade model validation reports. AI-specific governance documentation (Colorado AI Act, ISO 42001) trails platform certifications.	FICO Trust

DEEP DIVE

FICO is the default safe-choice AI vendor for US banks because the SR 11-7 documentation packet is already what every examiner expects. Forty-plus years of credit-scoring model validation is now extended to ML-driven fraud detection (Falcon) and credit scoring (FICO Score 10 T). The governance posture is the strongest in the banking category because validation isn't an add-on — it's the product.

STRENGTHS

- Reference SR 11-7 validation documentation
- FedRAMP + SOC 2 + ISO 27001 compliance stack
- BAA-eligible for PHI overlap; DPA standard for enterprise
- Public Responsible AI framework with NIST AI RMF mapping

WEAKNESSES

- No ISO/IEC 42001 attestation
- Pricing structure can be opaque at smaller community-bank scale
- AI-specific governance documentation trails core platform certifications

BEST USE CASE

Mid-market and large US banks running fraud detection or credit decisioning where examiner expectations have already standardized on FICO documentation. Lowest-friction SR 11-7 audit posture in the banking category.

AVOID WHEN

Smaller community banks (under $500M AUM) where the licensing economics don't amortize and lighter-weight alternatives like Hummingbird (AML) or Unit21 (transaction monitoring) match the actual exposure.

Last reviewed: 2026-05-13 · Homepage: https://www.fico.com · Trust center: https://www.fico.com/en/trust

12-AXIS SCORING

Axis	Status	Note	Source
BAA / DPA available	Yes	LexisNexis enterprise data-handling agreements address client-confidential data for firms.	LexisNexis Privacy
Training-data opt-out	Yes	Lexis+ AI does not train on customer prompts or content. Tenant isolation enforced.	LexisNexis Lexis+ AI Privacy
US data residency option	Yes	US data residency available for US customers; Lexis+ AI hosted on US infrastructure.	LexisNexis Trust
SOC 2 Type II report	Yes	LexisNexis platform holds SOC 2 Type II and ISO 27001.	LexisNexis Trust
ISO/IEC 42001 attestation	No	No ISO/IEC 42001 attestation as of May 2026.	Public posture review
NIST AI RMF self-attestation	Partial	LexisNexis publishes Responsible AI principles; no formal NIST AI RMF self-attestation.	LexisNexis Responsible AI
Colorado AI Act readiness	No	No Colorado AI Act-specific public statement.	Public posture review
HHS-OCR Section 1557 readiness	N/A	Legal-vertical positioning.	Lexis+ AI positioning
FRB SR 11-7 readiness	N/A	Legal-vertical positioning.	Lexis+ AI positioning
ABA Formal Op 512 readiness	Yes	LexisNexis publishes ABA Op 512 alignment documentation for Lexis+ AI.	LexisNexis Lexis+ AI ABA Op 512 documentation
Subprocessor list public	Yes	LexisNexis subprocessor list available via standard enterprise terms.	LexisNexis Subprocessors
Trust-center maturity	4/5	Mature LexisNexis platform compliance documentation. AI-specific governance present but less granular than cloud-platform peers.	LexisNexis Trust

DEEP DIVE

Lexis+ AI is the direct Lexis-content counterpart to CoCounsel. The governance posture is roughly equivalent on platform fundamentals (BAA, residency, SOC 2, ABA Op 512). The differentiator is which legal content corpus the firm has standardized on. Both are appropriate for ABA Op 512-aware deployment.

STRENGTHS

- Citation grounding from Lexis case-law and secondary sources
- ABA Op 512 alignment documented
- Default no-train, US residency, BAA-equivalent
- Inherits LexisNexis platform compliance stack

WEAKNESSES

- No ISO/IEC 42001
- No Colorado AI Act-specific public statement
- No formal NIST AI RMF self-attestation

BEST USE CASE

Firms standardized on Lexis content, where Lexis+ AI's content integration matches existing research workflows.

AVOID WHEN

Firms standardized on Westlaw — the content integration advantage shifts to CoCounsel.

Last reviewed: 2026-05-13 · Homepage: https://www.lexisnexis.com/en-us/products/lexis-plus-ai.page

12-AXIS SCORING

Axis	Status	Note	Source
BAA / DPA available	Yes	Covered under Thomson Reuters enterprise data-handling agreements.	Thomson Reuters Trust Center
Training-data opt-out	Yes	Westlaw Precision AI does not train on customer research queries or content. Tenant isolation enforced.	Thomson Reuters Privacy
US data residency option	Yes	US data residency available for US customers.	Thomson Reuters Trust Center
SOC 2 Type II report	Yes	Thomson Reuters Cloud Platform holds SOC 2 Type II and ISO 27001.	Thomson Reuters Trust Center
ISO/IEC 42001 attestation	No	No ISO/IEC 42001 attestation as of May 2026.	Public posture review
NIST AI RMF self-attestation	Partial	Thomson Reuters AI Principles framework; no formal NIST AI RMF self-attestation.	Thomson Reuters AI Principles
Colorado AI Act readiness	No	No Colorado AI Act-specific public statement.	Public posture review
HHS-OCR Section 1557 readiness	N/A	Legal-vertical positioning.	Westlaw Precision AI positioning
FRB SR 11-7 readiness	N/A	Legal-vertical positioning.	Westlaw Precision AI positioning
ABA Formal Op 512 readiness	Yes	Thomson Reuters publishes ABA Op 512 alignment documentation applicable to Westlaw Precision AI.	Thomson Reuters Westlaw Precision AI ABA Op 512 documentation
Subprocessor list public	Yes	Thomson Reuters subprocessor list published.	Thomson Reuters Subprocessors
Trust-center maturity	4/5	Same trust posture as CoCounsel — mature platform compliance, less granular AI-specific governance.	Thomson Reuters Trust Center

DEEP DIVE

Westlaw Precision AI is the AI-assisted research overlay on Westlaw — most directly comparable to Lexis+ AI's research workflow rather than CoCounsel's drafting workflow. The governance posture mirrors CoCounsel because both run on the same Thomson Reuters Cloud Platform.

STRENGTHS

- Citation grounding from Westlaw primary sources
- ABA Op 512 alignment documented
- Inherits Thomson Reuters Cloud Platform compliance

WEAKNESSES

- No ISO/IEC 42001
- No Colorado AI Act-specific public statement
- Pricing tied to Westlaw Precision tier — not a standalone purchase

BEST USE CASE

Firms standardized on Westlaw who want AI-assisted research without moving to CoCounsel's drafting workflow.

AVOID WHEN

Firms standardized on Lexis — the research-content advantage shifts to Lexis+ AI.

Last reviewed: 2026-05-13 · Homepage: https://legal.thomsonreuters.com/en/products/westlaw-precision

12-AXIS SCORING

Axis	Status	Note	Source
BAA / DPA available	Yes	BAA available under the standard Microsoft Online Services HIPAA BAA — covers Copilot for Microsoft 365 within the M365 commercial environment.	Microsoft HIPAA BAA + Trust Center
Training-data opt-out	Yes	Customer data is not used to train foundation models. M365 Copilot prompts and responses stay within the tenant boundary.	Microsoft Copilot Trust Center
US data residency option	Yes	M365 Copilot inherits M365 tenant data residency — US tenants stay in US datacenters by default. Advanced Data Residency add-on available.	Microsoft 365 Data Residency
SOC 2 Type II report	Yes	M365 commercial environment holds SOC 2 Type II, SOC 1 Type II, SOC 3, ISO 27001, ISO 27017, ISO 27018, FedRAMP High, IRAP, and others.	Microsoft Service Trust Portal
ISO/IEC 42001 attestation	Partial	Microsoft has announced ISO/IEC 42001 alignment work; certification scope public for Azure AI services. M365 Copilot scope confirmation pending.	Microsoft Responsible AI Standard
NIST AI RMF self-attestation	Partial	Microsoft publishes a Responsible AI Standard and Transparency Report mapped against NIST AI RMF functions. No formal self-attestation document.	Microsoft Responsible AI Transparency Report
Colorado AI Act readiness	Partial	Microsoft published a Colorado AI Act readiness statement framing M365 Copilot as a general-purpose AI tool with deployer responsibility for high-risk uses.	Microsoft AI law tracker
HHS-OCR Section 1557 readiness	Partial	BAA in place. Section 1557 compliance is deployer responsibility for clinical decision use; Microsoft documents the technical controls available.	Microsoft HIPAA documentation
FRB SR 11-7 readiness	Partial	Microsoft documents model risk management controls; SR 11-7 validation remains deployer responsibility.	Microsoft Financial Services compliance
ABA Formal Op 512 readiness	Partial	Microsoft publishes legal-sector AI guidance covering matter wall configuration in Copilot. ABA Op 512 obligations remain firm-level.	Microsoft Legal industry resources
Subprocessor list public	Yes	Microsoft Online Services subprocessor list public and granular.	Microsoft Service Trust Portal — Subprocessors
Trust-center maturity	5/5	Microsoft Service Trust Portal is the gold-standard reference — public certificate library, audit reports under NDA, granular subprocessor and residency documentation.	Microsoft Service Trust Portal

DEEP DIVE

M365 Copilot has the most complete governance posture in the productivity category. BAA, no-train, US residency, full SOC/ISO stack, public subprocessor list, and the most mature trust portal in the market. The risk is operational rather than vendor: matter-wall and DLP configuration in M365 is where firms fail Copilot governance, not the underlying BAA.

STRENGTHS

- BAA under standard Microsoft Online Services HIPAA BAA
- Default no-train, US residency, full compliance stack
- Most mature trust portal of any AI vendor
- Inherits enterprise-grade M365 identity and DLP controls

WEAKNESSES

- ISO 42001 certification scope not yet confirmed for Copilot
- Sector-specific readiness (Section 1557, SR 11-7, ABA Op 512) is deployer responsibility — Microsoft provides controls, not turnkey compliance
- Matter-wall and DLP configuration is non-trivial; many deployments fail at the configuration layer

BEST USE CASE

Organizations already standardized on Microsoft 365 commercial with mature DLP, Conditional Access, and SharePoint/OneDrive governance in place. Lowest-friction enterprise AI rollout in the regulated mid-market.

AVOID WHEN

Tenants without DLP, label, or Conditional Access maturity — Copilot inherits the existing access surface, so a tenant with weak governance becomes a worse tenant with Copilot.

Last reviewed: 2026-05-13 · Homepage: https://www.microsoft.com/en-us/microsoft-365/copilot · Trust center: https://servicetrust.microsoft.com

12-AXIS SCORING

Axis	Status	Note	Source
BAA / DPA available	Yes	Harvey signs enterprise data-handling agreements equivalent to BAA scope for client-confidential workloads. Firm-level deployment terms address privilege handling.	Harvey Security
Training-data opt-out	Yes	Harvey does not train on client data. Tenant isolation contractually enforced. Foundation models accessed via Harvey are configured with zero-retention enterprise contracts.	Harvey Security
US data residency option	Yes	US data residency available for enterprise customers. Harvey runs primarily on Azure US regions.	Harvey Security
SOC 2 Type II report	Yes	SOC 2 Type II completed. Report available to enterprise customers via direct request.	Harvey Security
ISO/IEC 42001 attestation	No	No public ISO/IEC 42001 attestation as of May 2026.	Public posture review
NIST AI RMF self-attestation	Partial	Harvey publishes governance documentation aligned to NIST AI RMF principles. No formal self-attestation.	Harvey governance documentation
Colorado AI Act readiness	Partial	Harvey acknowledges Colorado AI Act deployer responsibility model in customer documentation; firms own end-deployer obligations.	Harvey customer documentation
HHS-OCR Section 1557 readiness	N/A	Legal-vertical positioning.	Harvey positioning review
FRB SR 11-7 readiness	N/A	Legal-vertical positioning.	Harvey positioning review
ABA Formal Op 512 readiness	Yes	Harvey publishes ABA Formal Op 512 alignment documentation: data isolation, no training on client data, audit logging, privilege-aware retention controls.	Harvey ABA Op 512 documentation
Subprocessor list public	Partial	Subprocessor information available to enterprise customers under NDA. Not self-serve public.	Harvey enterprise documentation
Trust-center maturity	3/5	Security page documents core controls; enterprise-grade documentation available on request. Less self-serve maturity than cloud-platform vendors.	harvey.ai/security

DEEP DIVE

Harvey is the highest-profile legal vertical AI vendor. The governance posture is strong on the dimensions that matter most for law firms (no-train, US residency, BAA-equivalent, ABA Op 512 alignment) but trust-portal maturity lags cloud-platform vendors. The competitive position depends on the firm-specific workflow value rather than cross-cutting governance differentiation.

STRENGTHS

- Purpose-built for legal — privilege handling and matter walls native to product
- ABA Op 512 alignment documented
- Default no-train, US residency, BAA-equivalent
- Foundation-model upstreams contractually configured for zero-retention

WEAKNESSES

- No ISO/IEC 42001
- No formal NIST AI RMF self-attestation
- Trust portal less mature than cloud-platform peers
- Subprocessor transparency NDA-gated

BEST USE CASE

Am Law 100/200 firms with established AI governance, where Harvey's privilege-aware workflow and matter-context features deliver value beyond what a foundation model alone provides.

AVOID WHEN

Smaller firms (under 50 attorneys) where the per-attorney pricing doesn't amortize, and the ChatGPT Enterprise + ABA Op 512 protocol delivers acceptable functionality at lower cost.

Last reviewed: 2026-05-13 · Homepage: https://www.harvey.ai

12-AXIS SCORING

Axis	Status	Note	Source
BAA / DPA available	Yes	Zest AI signs DPAs / data-handling agreements for enterprise customers. BAA available where PHI exposure is in scope.	Zest AI Security
Training-data opt-out	Yes	Customer underwriting data not used for cross-customer model training. Tenant isolation enforced.	Zest AI Privacy
US data residency option	Yes	US data residency standard for US customers.	Zest AI Security
SOC 2 Type II report	Yes	Zest AI holds SOC 2 Type II.	Zest AI Security
ISO/IEC 42001 attestation	No	No ISO/IEC 42001 attestation as of May 2026.	Public posture review
NIST AI RMF self-attestation	Partial	Zest publishes Responsible AI documentation mapped to NIST AI RMF principles.	Zest AI Responsible AI
Colorado AI Act readiness	Partial	Zest has engaged on Colorado AI Act high-risk classification for credit decisioning.	Zest AI customer documentation
HHS-OCR Section 1557 readiness	N/A	Banking-vertical positioning.	Zest AI positioning
FRB SR 11-7 readiness	Yes	Zest publishes SR 11-7-grade model validation, ongoing monitoring, and fair-lending audit documentation. CFPB Circular 2023-03 adverse-action explainability built into the output format.	Zest AI SR 11-7 documentation
ABA Formal Op 512 readiness	N/A	Banking-vertical positioning.	Zest AI positioning
Subprocessor list public	Partial	Subprocessor list available to enterprise customers under NDA.	Zest AI Security
Trust-center maturity	3/5	Strong fair-lending + SR 11-7 documentation. Trust portal less self-serve than FICO; documentation distribution via enterprise relationship.	Zest AI Security

DEEP DIVE

Zest AI is the strongest pure-play banking AI vendor on fair-lending defensibility. The adverse-action explainability output is designed for CFPB Circular 2023-03 — explanations are model-derived rather than post-hoc, which matters in supervisory examination. Best fit for community and mid-size banks that need SR 11-7-aligned underwriting without standing up internal MRM capacity.

STRENGTHS

- CFPB Circular 2023-03 adverse-action explainability built into output
- SR 11-7-grade model validation documentation
- Tenant-isolated, US residency, BAA-eligible
- Purpose-built for fair-lending defensibility

WEAKNESSES

- No ISO/IEC 42001
- Trust portal less mature than FICO
- Smaller subprocessor transparency

BEST USE CASE

Community and mid-size banks ($500M-$10B AUM) deploying AI for personal lending, auto, or small-business decisioning where fair-lending audit defensibility is the binding constraint.

AVOID WHEN

Very large banks with deep internal MRM capacity may prefer to build on FICO or in-house given the volume.

Last reviewed: 2026-05-13 · Homepage: https://www.zest.ai

12-AXIS SCORING

Axis	Status	Note	Source
BAA / DPA available	Yes	Upstart signs DPAs and data-handling agreements with partner banks. BAA-eligible where PHI exposure exists in partner-bank datasets.	Upstart Security
Training-data opt-out	Yes	Partner-bank customer data processed under contracted purpose limitation. Cross-bank model training only with consortium consent.	Upstart Privacy
US data residency option	Yes	US data residency standard.	Upstart Security
SOC 2 Type II report	Yes	Upstart holds SOC 2 Type II.	Upstart Security
ISO/IEC 42001 attestation	No	No ISO/IEC 42001 attestation.	Public posture review
NIST AI RMF self-attestation	Partial	Upstart publishes Responsible AI + fair-lending governance documentation.	Upstart Responsible AI
Colorado AI Act readiness	Partial	Upstart has publicly engaged on Colorado AI Act readiness for credit decisioning.	Upstart customer documentation
HHS-OCR Section 1557 readiness	N/A	Banking-vertical positioning.	Upstart positioning
FRB SR 11-7 readiness	Yes	Upstart has CFPB no-action letter history (Sept 2017 + 2020 renewal) — uniquely deep fair-lending audit defensibility. SR 11-7-grade validation documentation maintained for partner-bank examiner needs.	CFPB No-Action Letter history
ABA Formal Op 512 readiness	N/A	Banking-vertical positioning.	Upstart positioning
Subprocessor list public	Partial	Subprocessor list available to enterprise customers.	Upstart Security
Trust-center maturity	3/5	Mature security documentation; CFPB engagement history is the differentiating compliance artifact. Trust portal less self-serve than enterprise platform vendors.	Upstart Security

DEEP DIVE

Upstart is uniquely defensible on fair-lending because of the CFPB no-action letter history — no other US AI lending vendor has that paper trail. The white-label partner model lets community banks deploy AI lending under Upstart's compliance umbrella, which is operationally easier than standing up internal validation. The cost is platform dependence: partner banks operate within Upstart's product roadmap rather than building proprietary capability.

STRENGTHS

- CFPB no-action letter history (Sept 2017 + 2020 renewal)
- Fair-lending audit defensibility uniquely deep
- Partner-bank model — origination under Upstart compliance umbrella
- SR 11-7-grade validation maintained for partner needs

WEAKNESSES

- Platform dependence — partner banks operate within Upstart's roadmap
- No ISO/IEC 42001
- Subprocessor transparency NDA-gated

BEST USE CASE

Community banks and credit unions wanting AI-driven personal lending or auto origination without internal model risk management capacity. The CFPB engagement history reduces partner-bank examiner risk.

AVOID WHEN

Banks that want proprietary AI capability or are concerned about platform dependence — building on FICO or licensing Zest AI keeps decisioning closer to in-house.

Last reviewed: 2026-05-13 · Homepage: https://www.upstart.com

12-AXIS SCORING

Axis	Status	Note	Source
BAA / DPA available	Yes	Suki signs BAAs for enterprise customers.	Suki Security
Training-data opt-out	Yes	Suki does not train models on customer audio or notes.	Suki Security
US data residency option	Yes	Suki US-hosted on US cloud infrastructure.	Suki Security
SOC 2 Type II report	Yes	Suki holds SOC 2 Type II and HITRUST CSF certification.	Suki Security
ISO/IEC 42001 attestation	No	No ISO/IEC 42001 attestation as of May 2026.	Public posture review
NIST AI RMF self-attestation	Partial	Suki publishes governance documentation aligning with NIST AI RMF principles; no formal self-attestation.	Suki Responsible AI
Colorado AI Act readiness	Partial	Suki engages on the Colorado AI Act deployer-responsibility model in customer documentation.	Suki customer documentation
HHS-OCR Section 1557 readiness	Partial	Suki documents bias testing and clinical safety governance; explicit Section 1557 public statement less detailed than Abridge.	Suki governance documentation
FRB SR 11-7 readiness	N/A	Healthcare-vertical positioning.	Suki positioning
ABA Formal Op 512 readiness	N/A	Healthcare-vertical positioning.	Suki positioning
Subprocessor list public	Yes	Subprocessor list available to enterprise customers.	Suki Security
Trust-center maturity	4/5	Mature security documentation with HITRUST + SOC 2. AI-specific governance less granular than Abridge.	Suki Security

DEEP DIVE

Suki has strong fundamentals — BAA, US residency, SOC 2, HITRUST — and a more pragmatic positioning than Abridge. The Section 1557 engagement is less prominent than Abridge but adequate for most ambulatory deployments. HITRUST CSF certification is a meaningful differentiator for health-system buyers that require it.

STRENGTHS

- BAA, US residency, SOC 2 Type II + HITRUST CSF
- Broad EHR integration
- Default no-train, customer-isolated

WEAKNESSES

- No ISO/IEC 42001
- Section 1557 documentation less prominent than Abridge
- Smaller scale than DAX Copilot or Abridge in market

BEST USE CASE

Ambulatory practices needing HITRUST-aligned procurement, broad EHR integration, and strong clinician workflow fit.

AVOID WHEN

Hospital systems with active OCR Section 1557 scrutiny — Abridge's public Section 1557 engagement is more defensible during audit.

Last reviewed: 2026-05-13 · Homepage: https://www.suki.ai · Trust center: https://www.suki.ai/security

12-AXIS SCORING

Axis	Status	Note	Source
BAA / DPA available	Yes	DAX Copilot is covered under Microsoft Online Services HIPAA BAA. Inherits the full M365/Azure BAA scope.	Microsoft Nuance DAX HIPAA
Training-data opt-out	Yes	Clinical encounter audio and generated notes are not used for foundation-model training. Customer-isolated processing.	Nuance DAX Copilot documentation
US data residency option	Yes	US data residency via Azure US regions. Customer-configurable.	Microsoft Azure Data Residency
SOC 2 Type II report	Yes	Microsoft Azure / M365 commercial environment compliance stack applies (SOC 2 Type II + SOC 1 + SOC 3 + ISO 27001/17/18 + FedRAMP).	Microsoft Service Trust Portal
ISO/IEC 42001 attestation	No	No DAX Copilot-specific ISO/IEC 42001 attestation as of May 2026.	Microsoft Service Trust Portal
NIST AI RMF self-attestation	Partial	Microsoft Responsible AI framework applies. No DAX-specific NIST AI RMF self-attestation document.	Microsoft Responsible AI
Colorado AI Act readiness	No	No DAX-specific Colorado AI Act public statement.	Public posture review
HHS-OCR Section 1557 readiness	Partial	BAA in place. Section 1557 algorithmic non-discrimination obligations for clinical decision support remain deployer responsibility; Microsoft documents the technical controls.	Microsoft Healthcare compliance
FRB SR 11-7 readiness	N/A	Healthcare-vertical positioning.	DAX positioning
ABA Formal Op 512 readiness	N/A	Healthcare-vertical positioning.	DAX positioning
Subprocessor list public	Yes	Microsoft Online Services subprocessor list applies.	Microsoft Service Trust Portal
Trust-center maturity	5/5	Inherits Microsoft Service Trust Portal — the gold-standard reference. DAX-specific documentation present on the Nuance side.	Microsoft Service Trust Portal

DEEP DIVE

DAX Copilot has the strongest healthcare-vertical governance posture in the market because it inherits the Microsoft/Azure/M365 compliance stack while being healthcare-positioned at the product layer. The result is best-in-class platform compliance combined with clinical workflow fit. The remaining gap is Section 1557 readiness, where the deployer still owns clinical-decision-support validation.

STRENGTHS

- Inherits Microsoft/Azure HIPAA BAA, US residency, SOC 2, ISO 27k, FedRAMP
- EHR-integrated (Epic, Cerner, athenahealth, etc.)
- Default no-train, customer-isolated processing
- Most mature trust portal of any healthcare AI vendor

WEAKNESSES

- No DAX-specific ISO/IEC 42001
- No Colorado AI Act-specific statement
- Section 1557 clinical-decision-support readiness is deployer-side

BEST USE CASE

Health systems and clinics with Microsoft 365 / Azure standardization where DAX Copilot's EHR integration matches the deployed EHR (Epic + DAX is the highest-leverage combination).

AVOID WHEN

Practices on EHRs without DAX integration (some smaller specialty EHRs) — the workflow value depends on EHR integration depth.

Last reviewed: 2026-05-13 · Homepage: https://www.nuance.com/healthcare/dragon-ai-clinical-solutions/dax-copilot.html

12-AXIS SCORING

Axis	Status	Note	Source
BAA / DPA available	Yes	BAA available under Salesforce Health Cloud and applicable to Einstein/Agentforce within the BAA-covered environment.	Salesforce HIPAA compliance
Training-data opt-out	Yes	Einstein Trust Layer enforces zero data retention by the underlying LLM provider. Customer data never used for model training.	Einstein Trust Layer
US data residency option	Yes	Salesforce supports US data residency through US-based Hyperforce regions. Customer-configurable.	Salesforce Hyperforce
SOC 2 Type II report	Yes	Salesforce holds SOC 2 Type II, SOC 1, ISO 27001/17/18, FedRAMP, and additional sector certifications.	Salesforce Compliance
ISO/IEC 42001 attestation	No	No ISO/IEC 42001 attestation for Einstein/Agentforce as of May 2026.	Salesforce Compliance
NIST AI RMF self-attestation	Partial	Salesforce publishes a Trusted AI Principles framework with explicit mapping to NIST AI RMF functions. No formal self-attestation document.	Salesforce Trusted AI
Colorado AI Act readiness	No	No Colorado AI Act-specific public statement; Salesforce documents the deployer responsibility model.	Public posture review
HHS-OCR Section 1557 readiness	Partial	BAA available; Section 1557 compliance for clinical decision support is deployer responsibility. Salesforce Health Cloud documents the technical controls.	Salesforce Health Cloud compliance
FRB SR 11-7 readiness	Partial	Salesforce Financial Services Cloud documents model risk controls; SR 11-7 validation is deployer responsibility.	Salesforce Financial Services compliance
ABA Formal Op 512 readiness	N/A	Not legal-vertical positioned.	Salesforce positioning review
Subprocessor list public	Yes	Salesforce subprocessor list public and granular.	Salesforce Subprocessors
Trust-center maturity	5/5	Mature compliance portal at compliance.salesforce.com — public certificates, subprocessor list, audit reports, sector-specific BAA addenda.	Salesforce Compliance

DEEP DIVE

Salesforce's governance posture is one of the strongest in the enterprise category because Einstein/Agentforce inherits the Salesforce platform compliance stack — BAA, US residency, FedRAMP, SOC 2, granular subprocessors. The Einstein Trust Layer's zero-retention enforcement at the LLM-provider boundary is operationally meaningful. The gap is sector-specific posture: deployers still own clinical or financial validation work.

STRENGTHS

- BAA, US residency, FedRAMP — full platform compliance stack
- Einstein Trust Layer enforces zero LLM-provider retention
- Most mature compliance portal in the productivity category
- Vertical Cloud (Health, Financial Services) integration

WEAKNESSES

- No ISO/IEC 42001
- No Colorado AI Act-specific statement
- Section 1557 / SR 11-7 readiness is deployer-side

BEST USE CASE

Salesforce-standardized organizations rolling out Agentforce within existing Health Cloud / Financial Services Cloud / Einstein Trust Layer configuration — governance inherits cleanly from the platform.

AVOID WHEN

Organizations without an existing Salesforce platform — the value of Einstein governance depends entirely on platform standardization.

Last reviewed: 2026-05-13 · Homepage: https://www.salesforce.com/products/einstein · Trust center: https://compliance.salesforce.com

12-AXIS SCORING

Axis	Status	Note	Source
BAA / DPA available	Yes	BAA available for enterprise customers. Glean supports HIPAA-covered deployments.	Glean Trust
Training-data opt-out	Yes	Customer data not used to train Glean's models. Default tenant isolation.	Glean Trust
US data residency option	Yes	US data residency option available for enterprise customers (US-only deployment).	Glean Trust
SOC 2 Type II report	Yes	SOC 2 Type II, ISO 27001:2022, ISO 27017, ISO 27018.	Glean Trust
ISO/IEC 42001 attestation	No	No ISO/IEC 42001 attestation as of May 2026.	Glean Trust
NIST AI RMF self-attestation	Partial	Public governance documentation aligns with NIST AI RMF functions; no formal self-attestation.	Glean Responsible AI
Colorado AI Act readiness	No	No Colorado AI Act-specific public statement.	Public posture review
HHS-OCR Section 1557 readiness	N/A	Not positioned for clinical decision support.	Glean positioning review
FRB SR 11-7 readiness	N/A	Not positioned as a banking decisioning system.	Glean positioning review
ABA Formal Op 512 readiness	N/A	Not legal-vertical positioned.	Glean positioning review
Subprocessor list public	Yes	Subprocessor list available to customers via the trust portal.	Glean Trust — Subprocessors
Trust-center maturity	4/5	Mature trust portal with public certificate library, audit reports under NDA, customer-facing documentation. Lacks AI-specific certifications (ISO 42001) and explicit Colorado AI Act statement.	Glean Trust

DEEP DIVE

Glean is an interesting governance case because it sits between cloud productivity tools and AI agents — permission-aware enterprise search that doesn't store source content but does perform retrieval-augmented generation. The governance stack is strong on the platform fundamentals (BAA, residency, SOC 2 + ISO) but doesn't claim sector-specific readiness because it's not a decisioning system.

STRENGTHS

- BAA + US residency + SOC 2 + ISO 27k stack
- Permission-aware retrieval respects source-system ACLs
- Default tenant isolation, no cross-customer training
- Mature subprocessor transparency

WEAKNESSES

- No ISO/IEC 42001
- No Colorado AI Act compliance statement
- Sector overlays (Section 1557, SR 11-7, ABA Op 512) not in scope by positioning

BEST USE CASE

Mid-market and enterprise organizations needing AI-grade enterprise search across a SaaS stack, with HIPAA BAA or general regulated-data handling requirements.

AVOID WHEN

Use cases that need vendor-side decisioning support — Glean is retrieval and answer-generation, not regulated-decision automation.

Last reviewed: 2026-05-13 · Homepage: https://www.glean.com · Trust center: https://www.glean.com/trust

12-AXIS SCORING

Axis	Status	Note	Source
BAA / DPA available	Yes	Arctic Wolf signs BAAs for healthcare customers handling PHI within scope of MDR telemetry.	Arctic Wolf Trust Center
Training-data opt-out	Yes	Customer telemetry is not used for cross-customer model training; tenant data remains in customer-scoped pipelines.	Arctic Wolf Trust Center
US data residency option	Yes	US data centers available; region configurable per customer engagement.	Arctic Wolf Trust Center
SOC 2 Type II report	Yes	SOC 2 Type II, ISO 27001, HIPAA, and PCI DSS attestations all held; reports available under NDA via Trust Center.	Arctic Wolf Trust Center
ISO/IEC 42001 attestation	No	No ISO/IEC 42001 AI management system attestation as of May 2026.	Public posture review
NIST AI RMF self-attestation	Partial	AI-augmented detection features documented in product materials but no formal NIST AI RMF self-attestation document published.	Arctic Wolf product documentation
Colorado AI Act readiness	No	No Colorado AI Act SB 24-205 readiness statement. MDR services are platform-neutral; downstream customer scope.	Public posture review
HHS-OCR Section 1557 readiness	N/A	MSSP — platform-neutral; Section 1557 algorithmic non-discrimination obligation sits with the healthcare customer.	Arctic Wolf positioning
FRB SR 11-7 readiness	N/A	MSSP — SR 11-7 model risk obligation sits with the financial institution customer.	Arctic Wolf positioning
ABA Formal Op 512 readiness	N/A	MSSP — ABA Formal Opinion 512 obligation sits with the law firm customer.	Arctic Wolf positioning
Subprocessor list public	Yes	Subprocessor list public via Trust Center.	Arctic Wolf Trust Center
Trust-center maturity	4/5	Mature trust center with SOC 2, ISO 27001, HIPAA, PCI documentation. AI-specific governance documentation lighter than platform compliance posture.	Arctic Wolf Trust Center

DEEP DIVE

Arctic Wolf's Concierge model with a named Concierge Security Team is the closest peer in the US MDR market to EFROS's named-senior-analyst positioning. Platform compliance is strong; AI features function as detection acceleration rather than autonomous response. The CST is the differentiator — customers get a named team rather than rotating tier-1 analysts.

STRENGTHS

- Named Concierge Security Team accountability model
- SOC 2 Type II + ISO 27001 + HIPAA + PCI all held
- US data residency standard with configurable region
- Subprocessor list published

WEAKNESSES

- No ISO/IEC 42001 AI management system attestation
- No Colorado AI Act readiness statement
- AI-specific governance documentation thinner than platform compliance
- Standard playbook constraints — customization beyond defaults is engagement-dependent

BEST USE CASE

Mid-market organizations wanting outsourced MDR with named-team accountability across endpoint, cloud, network, and identity, where the operational tempo of a standardized concierge playbook is a feature rather than a constraint.

AVOID WHEN

Customers needing deep customization or pre-authorized containment actions beyond Arctic Wolf's standard playbook, or environments requiring AI-decisioning transparency at the model level rather than detection-output level.

Last reviewed: 2026-05-13 · Homepage: https://arcticwolf.com · Trust center: https://arcticwolf.com/about-us/trust-center/

12-AXIS SCORING

Axis	Status	Note	Source
BAA / DPA available	Yes	Huntress signs BAAs for healthcare customers where PHI overlaps with telemetry scope.	Huntress Trust
Training-data opt-out	Yes	Customer telemetry not used for cross-customer model training; tenant data is scoped to the customer's environment.	Huntress Trust
US data residency option	Yes	US data residency standard.	Huntress Trust
SOC 2 Type II report	Yes	SOC 2 Type II report available via Trust portal; reports gated under NDA.	Huntress Trust
ISO/IEC 42001 attestation	No	No ISO/IEC 42001 attestation.	Public posture review
NIST AI RMF self-attestation	Partial	AI-augmented threat hunting features documented; no formal NIST AI RMF self-attestation document.	Huntress product documentation
Colorado AI Act readiness	No	No Colorado AI Act readiness statement.	Public posture review
HHS-OCR Section 1557 readiness	N/A	MSSP — Section 1557 obligation sits with the healthcare customer.	Huntress positioning
FRB SR 11-7 readiness	N/A	MSSP — SR 11-7 obligation sits with the financial institution customer.	Huntress positioning
ABA Formal Op 512 readiness	N/A	MSSP — ABA Formal Opinion 512 obligation sits with the law firm customer.	Huntress positioning
Subprocessor list public	Yes	Subprocessor list public via Trust portal.	Huntress Trust
Trust-center maturity	4/5	Trust portal includes SOC 2, subprocessor list, security documentation. AI governance documentation lighter than platform compliance posture.	Huntress Trust

DEEP DIVE

Huntress is best-in-class for endpoint and M365 identity threat detection at the SMB-to-mid-market scale. The AI features function as decision-support for human threat hunters rather than autonomous response. Distribution is partner-led (MSP channel + direct), and pricing is calibrated below enterprise MDR.

STRENGTHS

- Strong endpoint and M365 identity coverage for the price point
- SOC 2 Type II, US residency, BAA available
- Subprocessor transparency via Trust portal
- Decision-support AI keeps human-in-the-loop accountability clear

WEAKNESSES

- No ISO/IEC 42001 attestation
- No Colorado AI Act readiness statement
- Coverage scope intentionally narrower than full-XDR MDR (no native network or OT)
- AI-specific governance documentation thinner than platform compliance

BEST USE CASE

Organizations with limited internal security capacity wanting strong endpoint and M365 identity threat detection without paying enterprise MDR pricing. Particularly strong fit for MSP-distributed delivery to SMB end customers.

AVOID WHEN

Enterprises needing full-spectrum XDR with native network, OT, or cloud workload protection — Huntress's coverage is intentionally focused rather than comprehensive.

Last reviewed: 2026-05-13 · Homepage: https://www.huntress.com · Trust center: https://www.huntress.com/trust

12-AXIS SCORING

Axis	Status	Note	Source
BAA / DPA available	Yes	eSentire signs BAAs for healthcare customers; PHI scope addressed within MDR engagement.	eSentire Trust Center
Training-data opt-out	Yes	Customer telemetry not used for cross-customer model training within Atlas AI; tenant-scoped pipelines.	eSentire Trust Center
US data residency option	Yes	US data residency available; multi-region architecture with customer configuration.	eSentire Trust Center
SOC 2 Type II report	Yes	SOC 2 Type II, ISO 27001, HIPAA, PCI, and FedRAMP-aligned posture documented via Trust Center.	eSentire Trust Center
ISO/IEC 42001 attestation	No	No ISO/IEC 42001 attestation for the Atlas AI platform as of May 2026.	Public posture review
NIST AI RMF self-attestation	Partial	Atlas AI platform documented with model governance materials but no formal NIST AI RMF self-attestation published.	eSentire Atlas AI documentation
Colorado AI Act readiness	No	No Colorado AI Act SB 24-205 readiness statement.	Public posture review
HHS-OCR Section 1557 readiness	N/A	MSSP — Section 1557 obligation sits with the healthcare customer.	eSentire positioning
FRB SR 11-7 readiness	N/A	MSSP — SR 11-7 obligation sits with the financial institution customer.	eSentire positioning
ABA Formal Op 512 readiness	N/A	MSSP — ABA Formal Opinion 512 obligation sits with the law firm customer.	eSentire positioning
Subprocessor list public	Yes	Subprocessor list public via Trust Center.	eSentire Trust Center
Trust-center maturity	4/5	Mature trust center with full attestation stack and FedRAMP-aligned posture. Atlas AI platform branding is the most explicit AI-MDR positioning in the category, though formal AI governance attestation (ISO 42001) is absent.	eSentire Trust Center

DEEP DIVE

eSentire's Atlas AI is the most explicit AI-platform branding in the MDR category and threat hunt depth is the operational differentiator. The TRU (Threat Response Unit) does proprietary detection engineering paired with AI augmentation. Best fit for enterprises that prioritize hunt depth over coverage breadth.

STRENGTHS

- Full attestation stack — SOC 2, ISO 27001, HIPAA, PCI, FedRAMP-aligned
- Atlas AI platform with explicit AI-MDR positioning
- Threat Response Unit (TRU) proprietary detection engineering
- Subprocessor transparency via Trust Center

WEAKNESSES

- No ISO/IEC 42001 attestation for Atlas AI
- No Colorado AI Act readiness statement
- Premium pricing tier vs. SMB-focused MDR alternatives
- AI governance posture lighter than platform compliance maturity

BEST USE CASE

Enterprises that prioritize threat hunt depth over breadth — particularly those needing proprietary detection engineering against targeted threat actors rather than commodity malware coverage.

AVOID WHEN

Cost-sensitive SMBs where Huntress-tier coverage is sufficient, or organizations that need explicit ISO 42001 AI governance attestation as a procurement requirement.

Last reviewed: 2026-05-13 · Homepage: https://www.esentire.com · Trust center: https://www.esentire.com/about-us/trust-center

12-AXIS SCORING

Axis	Status	Note	Source
BAA / DPA available	Yes	Sophos signs BAAs for healthcare customers within scope of platform and MDR engagement.	Sophos Trust Center
Training-data opt-out	Yes	Customer data not used for cross-customer model training; Intercept X models updated via Sophos research pipeline rather than tenant data.	Sophos Trust Center
US data residency option	Yes	US data residency available via Sophos Central region configuration.	Sophos Trust Center
SOC 2 Type II report	Yes	SOC 2 and ISO 27001 held; reports available under NDA via Trust Center.	Sophos Trust Center
ISO/IEC 42001 attestation	No	No ISO/IEC 42001 attestation for Intercept X or Sophos AI features as of May 2026.	Public posture review
NIST AI RMF self-attestation	Partial	Sophos AI research publications and product documentation cover model governance themes; no formal NIST AI RMF self-attestation document published.	Sophos AI research
Colorado AI Act readiness	No	No Colorado AI Act readiness statement.	Public posture review
HHS-OCR Section 1557 readiness	N/A	MSSP / platform vendor — Section 1557 obligation sits with the healthcare customer.	Sophos positioning
FRB SR 11-7 readiness	N/A	MSSP / platform vendor — SR 11-7 obligation sits with the financial institution customer.	Sophos positioning
ABA Formal Op 512 readiness	N/A	MSSP / platform vendor — ABA Formal Opinion 512 obligation sits with the law firm customer.	Sophos positioning
Subprocessor list public	Yes	Subprocessor list public via Trust Center.	Sophos Trust Center
Trust-center maturity	4/5	Mature trust center with SOC 2, ISO 27001, subprocessor list, and active AI research publications. AI governance documentation is product-research-led rather than formal attestation.	Sophos Trust Center

DEEP DIVE

Sophos AI is the longest-established AI in endpoint security — the Invincea acquisition in 2017 brought deep-learning malware detection into Intercept X well before the category was crowded. Sophos MDR overlays managed detection on top of the platform. Best fit for organizations wanting vendor-integrated endpoint AI without a separate MDR contract.

STRENGTHS

- Longest-running deep-learning endpoint AI lineage in the category
- SOC 2 + ISO 27001 + BAA + US residency standard
- Vendor-integrated stack — endpoint, firewall, MDR from one platform
- Active AI research publications

WEAKNESSES

- No ISO/IEC 42001 attestation
- No Colorado AI Act readiness statement
- Coverage breadth concentrated on endpoint + network — XDR depth varies by module
- AI governance documentation product-research-led rather than formal attestation

BEST USE CASE

Organizations wanting vendor-integrated endpoint AI without a separate MDR contract — particularly mid-market buyers who value a single-pane Sophos Central platform across endpoint, firewall, and managed detection.

AVOID WHEN

Enterprises needing full-spectrum XDR coverage beyond endpoint and network — cloud workload protection and identity threat detection are stronger in dedicated MDR competitors.

Last reviewed: 2026-05-13 · Homepage: https://www.sophos.com · Trust center: https://www.sophos.com/en-us/legal/trust-center

12-AXIS SCORING

Axis	Status	Note	Source
BAA / DPA available	Yes	Unit21 signs DPAs for enterprise customers; BAA available where PHI overlap exists.	Unit21 Security
Training-data opt-out	Yes	Customer transaction data not used for cross-customer model training.	Unit21 Privacy
US data residency option	Yes	US data residency standard.	Unit21 Security
SOC 2 Type II report	Yes	Unit21 holds SOC 2 Type II.	Unit21 Security
ISO/IEC 42001 attestation	No	No ISO/IEC 42001 attestation.	Public posture review
NIST AI RMF self-attestation	Partial	Unit21 publishes governance documentation aligned to NIST AI RMF; no formal self-attestation.	Unit21 Responsible AI
Colorado AI Act readiness	No	No Colorado AI Act-specific public statement.	Public posture review
HHS-OCR Section 1557 readiness	N/A	Banking-vertical positioning.	Unit21 positioning
FRB SR 11-7 readiness	Partial	Unit21 documents SR 11-7 model risk practices for partner banks; full validation packet typically delivered under enterprise engagement rather than self-serve.	Unit21 customer documentation
ABA Formal Op 512 readiness	N/A	Banking-vertical positioning.	Unit21 positioning
Subprocessor list public	Yes	Subprocessor list public via trust documentation.	Unit21 Security
Trust-center maturity	4/5	Mature security documentation, modern compliance stack, public subprocessor list. AI-specific governance documentation present but lighter than FICO/Zest.	Unit21 Security

DEEP DIVE

Unit21 is the modern transaction-monitoring + fraud detection platform built for fintech-era institutions. The governance posture is solid on platform fundamentals (SOC 2, DPA, US residency, subprocessor transparency) and improving on AI-specific governance — but trails the pure-play SR 11-7 vendors (FICO, Zest) on validation packet depth. Best fit for institutions whose legacy AML vendor doesn't match their operational model.

STRENGTHS

- SOC 2 Type II, US residency, DPA standard
- Modern transaction-monitoring architecture
- Public subprocessor list
- Default tenant isolation

WEAKNESSES

- No ISO/IEC 42001
- No Colorado AI Act statement
- SR 11-7 validation packet depth lighter than FICO/Zest

BEST USE CASE

Neobanks, payments processors, crypto-adjacent institutions, and fintech-aligned community banks where legacy AML/transaction-monitoring vendors don't fit the data model or operational tempo.

AVOID WHEN

Traditional banks where examiners already standardized on FICO Falcon or NICE Actimize — the migration cost may exceed the operational benefit.

Last reviewed: 2026-05-13 · Homepage: https://www.unit21.ai

12-AXIS SCORING

Axis	Status	Note	Source
BAA / DPA available	Yes	Ironclad signs BAAs for enterprise customers with PHI obligations.	Ironclad Trust
Training-data opt-out	Yes	Customer contract content not used for training Ironclad's AI models.	Ironclad Trust
US data residency option	Yes	US data residency available for enterprise customers.	Ironclad Trust
SOC 2 Type II report	Yes	Ironclad holds SOC 2 Type II, ISO 27001, ISO 27017, ISO 27018.	Ironclad Trust
ISO/IEC 42001 attestation	No	No ISO/IEC 42001 attestation.	Public posture review
NIST AI RMF self-attestation	No	No public NIST AI RMF self-attestation.	Public posture review
Colorado AI Act readiness	No	No Colorado AI Act-specific public statement.	Public posture review
HHS-OCR Section 1557 readiness	N/A	Not positioned for clinical use.	Ironclad positioning
FRB SR 11-7 readiness	N/A	Not positioned as a banking decisioning system.	Ironclad positioning
ABA Formal Op 512 readiness	Partial	Ironclad publishes general AI governance documentation; explicit ABA Op 512 mapping less prominent than legal-research-focused vendors.	Ironclad AI governance documentation
Subprocessor list public	Yes	Subprocessor list public via trust portal.	Ironclad Trust
Trust-center maturity	4/5	Mature trust portal with public certificate library, audit reports under NDA, subprocessor list. AI-specific governance less prominent than platform fundamentals.	ironcladapp.com/trust

DEEP DIVE

Ironclad is best understood as a CLM platform with AI features rather than a pure legal AI vendor. The governance posture is strong on platform fundamentals (BAA, residency, SOC 2 + ISO stack) — matches the standard a corporate legal team would require for any CLM. AI-specific governance is less prominent because the AI is an overlay on the contract workflow.

STRENGTHS

- BAA + US residency + SOC 2 + ISO 27k stack
- Mature trust portal
- Default no-train
- Public subprocessor list

WEAKNESSES

- No ISO/IEC 42001
- No NIST AI RMF self-attestation
- ABA Op 512 mapping less prominent than research-focused legal vendors

BEST USE CASE

In-house legal teams using Ironclad as primary CLM, where AI features are workflow overlays rather than standalone deliverables.

AVOID WHEN

Litigation or research-heavy practices — Ironclad's AI is contract-workflow-oriented, not research or matter-aware drafting.

Last reviewed: 2026-05-13 · Homepage: https://ironcladapp.com · Trust center: https://ironcladapp.com/trust

12-AXIS SCORING

Axis	Status	Note	Source
BAA / DPA available	Partial	BAA available for Claude for Work Enterprise and Anthropic API on opt-in. Free and Pro tiers have no BAA.	Anthropic Trust Center — HIPAA
Training-data opt-out	Yes	Default no-train across all paid tiers and the API. Free/Pro consumer prompts also not used for training by default since 2024.	Anthropic Privacy Policy
US data residency option	Partial	Hosted on AWS US-East. No documented residency configuration option for enterprise customers as of May 2026.	Anthropic Trust Center
SOC 2 Type II report	Yes	SOC 2 Type II report available through the Anthropic Trust Center under NDA. ISO 27001:2022 also held.	Anthropic Trust Center
ISO/IEC 42001 attestation	No	No ISO/IEC 42001 attestation as of May 2026.	Anthropic Trust Center certificate list
NIST AI RMF self-attestation	Partial	Public alignment through Anthropic's Responsible Scaling Policy and Acceptable Use Policy. No formal NIST AI RMF self-attestation.	Anthropic Responsible Scaling Policy
Colorado AI Act readiness	No	No public Colorado AI Act SB 24-205 compliance statement.	Public posture review
HHS-OCR Section 1557 readiness	N/A	Foundation model — downstream healthcare deployer owns Section 1557 obligation.	HHS-OCR Section 1557 — deployer scope
FRB SR 11-7 readiness	N/A	Foundation model — downstream financial institution owns SR 11-7 validation.	FRB SR 11-7 — deployer scope
ABA Formal Op 512 readiness	N/A	Foundation model — downstream law firm owns ABA Formal Opinion 512 obligation.	ABA Formal Op 512 — practitioner scope
Subprocessor list public	Yes	Subprocessor list public via trust center (AWS, Google Cloud, billing/payments processors).	Anthropic Trust Center — Subprocessors
Trust-center maturity	4/5	Active trust center with NDA-gated audit reports, public Responsible Scaling Policy and Usage Policy. No public ISO 42001 or Colorado AI Act statement.	Anthropic Trust Center

DEEP DIVE

Anthropic's posture is closest peer to OpenAI on enterprise governance. The differentiator is the explicit safety-research orientation — Constitutional AI, Responsible Scaling Policy, public model behavior commitments. Default no-train across all tiers is a meaningful win versus OpenAI's opt-out-required consumer tiers. Residency configurability is weaker than OpenAI.

STRENGTHS

- Default no-train across all tiers, including consumer
- BAA available for Claude for Work Enterprise + API
- Responsible Scaling Policy is the most explicit public AI safety commitment of any foundation vendor
- SOC 2 Type II + ISO 27001

WEAKNESSES

- No US data residency configuration option
- No ISO/IEC 42001
- No Colorado AI Act compliance statement
- BAA only on Enterprise + API — shadow-AI risk on Pro/Free tiers

BEST USE CASE

Regulated organizations adopting Claude for Work Enterprise with the BAA, where default no-train across all tiers reduces the consumer-tier leakage risk. Strongest fit for organizations where the Responsible Scaling Policy aligns with internal AI safety governance.

AVOID WHEN

Strict US-data-residency requirements where the contract calls for documented residency control (Anthropic has less mature residency configurability than OpenAI Enterprise).

Last reviewed: 2026-05-13 · Homepage: https://www.anthropic.com · Trust center: https://trust.anthropic.com

12-AXIS SCORING

Axis	Status	Note	Source
BAA / DPA available	Partial	BAA available for Gemini for Workspace and Vertex AI when covered under the existing Google Workspace BAA. Consumer Gemini at gemini.google.com is not BAA-covered.	Google Cloud HIPAA Compliance
Training-data opt-out	Partial	Workspace and Vertex AI inputs not used to train consumer models. Consumer Gemini conversations are stored and may be reviewed for product improvement unless manually disabled.	Google Gemini Apps Privacy
US data residency option	Yes	Vertex AI and Workspace support US data residency through Google Cloud regions. Documented configuration option.	Google Cloud Data Residency
SOC 2 Type II report	Yes	Google Cloud holds SOC 2 Type II, SOC 3, ISO 27001/17/18. Reports available through Compliance Reports Manager.	Google Cloud Compliance
ISO/IEC 42001 attestation	No	No ISO/IEC 42001 attestation for Gemini/Vertex AI as of May 2026.	Google Cloud Compliance
NIST AI RMF self-attestation	Partial	Public mapping through Google's AI Principles and the Google Cloud Secure AI Framework (SAIF). No formal NIST AI RMF self-attestation document.	Google Secure AI Framework
Colorado AI Act readiness	No	No public Colorado AI Act compliance statement for Gemini.	Public posture review
HHS-OCR Section 1557 readiness	N/A	Foundation model — downstream healthcare deployer owns Section 1557 obligation. (Med-PaLM is a separate offering with distinct posture.)	HHS-OCR Section 1557 — deployer scope
FRB SR 11-7 readiness	N/A	Foundation model — downstream financial institution owns SR 11-7 validation.	FRB SR 11-7 — deployer scope
ABA Formal Op 512 readiness	N/A	Foundation model — downstream law firm owns ABA Formal Opinion 512 obligation.	ABA Formal Op 512 — practitioner scope
Subprocessor list public	Yes	Google Cloud subprocessor list public and granular.	Google Cloud Subprocessors
Trust-center maturity	4/5	Mature Google Cloud trust center, broad compliance coverage. Loses a point because Gemini-specific AI governance documentation (Colorado AI Act, ISO 42001) lags behind cloud-side posture.	Google Cloud Trust Center

DEEP DIVE

Gemini's governance posture inherits from Google Cloud — strong on certifications, US residency, subprocessor transparency, BAA coverage. AI-specific governance (Colorado AI Act, ISO 42001) lags behind cloud-side maturity. The strongest fit is Workspace-standardized organizations where Gemini is a configuration toggle rather than a new vendor relationship.

STRENGTHS

- BAA via Google Workspace inheritance
- Mature US data residency via Vertex AI / Workspace
- Strong subprocessor transparency
- Cloud-side SOC 2 + ISO 27k coverage

WEAKNESSES

- No ISO/IEC 42001 attestation
- No Colorado AI Act compliance statement
- Consumer Gemini has weaker default privacy posture
- AI-governance documentation behind cloud-side maturity

BEST USE CASE

Workspace-standardized organizations that already have a Google Workspace BAA and US data-residency settings configured — Gemini deployment is a contract-line-item exercise rather than a new vendor onboarding.

AVOID WHEN

Organizations without Google Workspace standardization — the cloud-side posture is what makes Gemini governance work, and bolting it onto a non-Google environment loses most of the advantage.

Last reviewed: 2026-05-13 · Homepage: https://gemini.google.com · Trust center: https://cloud.google.com/trust-center

12-AXIS SCORING

Axis	Status	Note	Source
BAA / DPA available	Yes	Hummingbird signs DPAs for enterprise customers; BAA-eligible where PHI overlap exists.	Hummingbird Security
Training-data opt-out	Yes	Customer case data not used for cross-customer model training.	Hummingbird Privacy
US data residency option	Yes	US data residency standard.	Hummingbird Security
SOC 2 Type II report	Yes	Hummingbird holds SOC 2 Type II.	Hummingbird Security
ISO/IEC 42001 attestation	No	No ISO/IEC 42001 attestation.	Public posture review
NIST AI RMF self-attestation	No	No public NIST AI RMF self-attestation. Hummingbird positions primarily as a workflow tool rather than an AI decisioning system; AI features (investigation summarization, transaction analytics) score lighter on RMF posture.	Public posture review
Colorado AI Act readiness	No	No Colorado AI Act-specific public statement.	Public posture review
HHS-OCR Section 1557 readiness	N/A	Banking-vertical positioning.	Hummingbird positioning
FRB SR 11-7 readiness	Partial	Hummingbird workflow does not directly perform credit decisioning; SR 11-7 applies to upstream transaction-monitoring model vendors. Hummingbird documents the audit trail expected for examiner-facing case management.	Hummingbird customer documentation
ABA Formal Op 512 readiness	N/A	Banking-vertical positioning.	Hummingbird positioning
Subprocessor list public	Partial	Subprocessor list available to enterprise customers.	Hummingbird Security
Trust-center maturity	3/5	Security documentation mature; AI-specific governance documentation absent. Strong workflow audit-trail features for BSA/AML examiner readiness.	Hummingbird Security

DEEP DIVE

Hummingbird is best understood as an AML workflow + audit-trail platform with AI overlay, rather than a decisioning AI vendor. The governance posture reflects this — strong on platform fundamentals (SOC 2, DPA, US residency) but light on AI-specific governance (NIST AI RMF, Colorado AI Act). SR 11-7 applies indirectly: Hummingbird documents the workflow, but upstream transaction-monitoring vendors own model risk.

STRENGTHS

- SOC 2 Type II, US residency, DPA standard
- Mature BSA/AML workflow + examiner audit trail
- Default tenant isolation

WEAKNESSES

- No NIST AI RMF self-attestation
- No Colorado AI Act statement
- AI-specific governance documentation thin
- Workflow-positioned rather than AI decisioning — model risk lives upstream

BEST USE CASE

Community banks, credit unions, and crypto-adjacent institutions needing modern BSA/AML case management with examiner-ready audit trails. Pair with a dedicated transaction-monitoring model vendor (Unit21, Verafin, NICE Actimize) for the AI model risk piece.

AVOID WHEN

Institutions looking for a single-vendor BSA/AML AI solution — Hummingbird is workflow + investigation, not the underlying decisioning model.

Last reviewed: 2026-05-13 · Homepage: https://www.hummingbird.co

12-AXIS SCORING

Axis	Status	Note	Source
BAA / DPA available	Partial	BAA available for ChatGPT Enterprise and OpenAI API on opt-in. ChatGPT Free, Plus, and Team have no BAA — never use for PHI.	OpenAI Enterprise Privacy
Training-data opt-out	Partial	Enterprise/Team/API default to no-train on customer data. ChatGPT Plus and Free require manual opt-out via settings (data still used for safety/abuse monitoring).	OpenAI Data Controls FAQ
US data residency option	Partial	Data Residency in the US available for ChatGPT Enterprise/Edu and API. Not default — must be configured.	OpenAI Data Residency announcement
SOC 2 Type II report	Yes	SOC 2 Type II report available through OpenAI Trust Portal under NDA. ISO 27001:2022, 27017, 27018 also held.	OpenAI Trust Portal
ISO/IEC 42001 attestation	No	No ISO/IEC 42001 attestation as of May 2026. OpenAI publishes a Preparedness Framework and Model Spec but no third-party AI MS audit.	OpenAI Trust Portal certificate index
NIST AI RMF self-attestation	Partial	Public alignment via OpenAI's Preparedness Framework and Model Spec. No formal NIST AI RMF self-attestation document.	OpenAI Preparedness Framework
Colorado AI Act readiness	No	No public Colorado AI Act SB 24-205 compliance statement. Downstream deployers using OpenAI in high-risk decisions carry the compliance burden.	Public posture review
HHS-OCR Section 1557 readiness	N/A	Foundation model — downstream healthcare deployer owns Section 1557 algorithmic non-discrimination obligation.	HHS-OCR Section 1557 Final Rule (May 2024) — deployer scope
FRB SR 11-7 readiness	N/A	Foundation model — downstream financial institution owns SR 11-7 validation responsibility.	FRB SR 11-7 — deployer scope
ABA Formal Op 512 readiness	N/A	Foundation model — downstream law firm owns ABA Formal Opinion 512 obligation.	ABA Formal Op 512 — practitioner scope
Subprocessor list public	Yes	Subprocessor list public (Microsoft Azure hosting, Stripe billing, Snowflake analytics, etc.).	OpenAI Enterprise Privacy — Subprocessors
Trust-center maturity	4/5	Active trust portal at trust.openai.com — audit reports under NDA, security whitepaper, public policy documents. Falls short of a 5 because no public ISO 42001 or Colorado AI Act statement yet.	OpenAI Trust Portal

DEEP DIVE

OpenAI is the highest-volume US AI vendor in regulated buyer pipelines. The governance posture is strong on the enterprise tier (BAA, no-train default, US data residency, SOC 2 + ISO 27k stack) and weak on consumer (no BAA, manual opt-out, no residency control). The single biggest deployment risk we see is staff using consumer ChatGPT for work where Enterprise was assumed.

STRENGTHS

- BAA available for ChatGPT Enterprise + API
- Default no-train on customer data at Enterprise/Team/API tiers
- Mature trust portal with under-NDA audit reports
- US data residency option for enterprise customers

WEAKNESSES

- No BAA on Plus/Team/Free — common shadow-AI source
- No ISO/IEC 42001 attestation as of May 2026
- No public Colorado AI Act compliance statement
- Sector-specific readiness (Section 1557, SR 11-7, ABA Op 512) is deployer responsibility — no vendor-side support

BEST USE CASE

Regulated organizations that have already standardized on ChatGPT Enterprise with the BAA in place, training opt-out enforced, and Data Residency in the US enabled — and have eliminated shadow consumer-tier use through DLP + identity policy.

AVOID WHEN

PHI workflows on ChatGPT Plus, Team, or Free; clinical decision support without a separately validated Section 1557 layer; bank credit decisioning without an SR 11-7 wrapper on top.

Last reviewed: 2026-05-13 · Homepage: https://openai.com · Trust center: https://trust.openai.com

12-AXIS SCORING

Axis	Status	Note	Source
BAA / DPA available	Partial	ConnectWise signs DPAs for the platform itself; BAA chain depends on the MSP's downstream contractual posture with end customers handling PHI.	ConnectWise Trust
Training-data opt-out	Yes	Customer data not used for cross-customer model training within Asio AI features.	ConnectWise Trust
US data residency option	Partial	Multi-region architecture; US residency available with customer configuration but not the default across all Asio modules.	ConnectWise Trust
SOC 2 Type II report	Yes	SOC 2 Type II held across core Asio platform modules.	ConnectWise Trust
ISO/IEC 42001 attestation	No	No ISO/IEC 42001 attestation.	Public posture review
NIST AI RMF self-attestation	No	No public NIST AI RMF self-attestation for Asio AI features as of May 2026.	Public posture review
Colorado AI Act readiness	No	No Colorado AI Act readiness statement.	Public posture review
HHS-OCR Section 1557 readiness	N/A	MSP platform — Section 1557 obligation sits with the downstream healthcare end customer, with the MSP as intermediate operator.	ConnectWise positioning
FRB SR 11-7 readiness	N/A	MSP platform — SR 11-7 obligation sits with the financial institution end customer.	ConnectWise positioning
ABA Formal Op 512 readiness	N/A	MSP platform — ABA Formal Opinion 512 obligation sits with the law firm end customer.	ConnectWise positioning
Subprocessor list public	Yes	Subprocessor list published.	ConnectWise Trust
Trust-center maturity	3/5	Platform compliance documentation is solid (SOC 2, subprocessor list) but AI-specific governance documentation is materially thinner than direct-to-enterprise MDR vendors. Distribution model is MSP-channel — governance posture reflects that downstream chain.	ConnectWise Trust

DEEP DIVE

ConnectWise is platform-and-channel rather than direct-to-enterprise — sold to MSPs who deliver downstream IT services. AI features in Asio accelerate MSP workflow (ticket automation, asset insights, PSA workflows) but the governance posture reflects the indirect distribution model. Platform fundamentals are solid; AI-specific documentation lags direct-MDR vendors.

STRENGTHS

- SOC 2 Type II across core Asio modules
- Public subprocessor list
- Training opt-out standard for Asio AI features
- Mature MSP-channel distribution and partner enablement

WEAKNESSES

- No NIST AI RMF self-attestation
- No ISO/IEC 42001 attestation
- No Colorado AI Act readiness statement
- BAA chain depends on downstream MSP contracts — not a single-vendor compliance answer for end customers

BEST USE CASE

MSPs delivering managed IT services to SMB and mid-market end customers, where AI features are workflow acceleration for the MSP operator rather than autonomous decisioning for end customers.

AVOID WHEN

Enterprises buying direct — ConnectWise's distribution model is MSP-channel, and the governance posture reflects that. Direct-to-enterprise MDR vendors are a closer match for direct buyers.

Last reviewed: 2026-05-13 · Homepage: https://www.connectwise.com

12-AXIS SCORING

Axis	Status	Note	Source
BAA / DPA available	Yes	Spellbook signs BAAs for enterprise customers where required.	Spellbook Security
Training-data opt-out	Yes	Spellbook does not train on customer documents. Tenant isolation enforced.	Spellbook Privacy
US data residency option	Partial	Spellbook hosted on US/Canada cloud infrastructure. Explicit US-only residency configuration not documented as of May 2026.	Spellbook Security
SOC 2 Type II report	Partial	Spellbook is SOC 2 Type II under audit / completed; report distribution via direct enterprise request.	Spellbook Security
ISO/IEC 42001 attestation	No	No ISO/IEC 42001 attestation.	Public posture review
NIST AI RMF self-attestation	No	No public NIST AI RMF self-attestation.	Public posture review
Colorado AI Act readiness	No	No Colorado AI Act-specific public statement.	Public posture review
HHS-OCR Section 1557 readiness	N/A	Legal-vertical positioning.	Spellbook positioning
FRB SR 11-7 readiness	N/A	Legal-vertical positioning.	Spellbook positioning
ABA Formal Op 512 readiness	Partial	Spellbook publishes general legal-ethics alignment documentation; explicit ABA Op 512 mapping less detailed than top-tier legal-vertical vendors.	Spellbook documentation
Subprocessor list public	Partial	Subprocessor information available via enterprise request; not self-serve public.	Spellbook Security
Trust-center maturity	2/5	Security page documents core controls. Trust-portal maturity below cloud-platform and top-tier legal-vertical peers.	spellbook.legal/security

DEEP DIVE

Spellbook targets a smaller-firm market than Harvey, Lexis+ AI, or CoCounsel. The governance posture reflects the smaller-vendor scale — solid fundamentals on the dimensions that matter most for contracts (BAA, no-train) but less mature on trust-portal documentation, sector-specific governance, and AI-specific certifications.

STRENGTHS

- BAA-eligible for enterprise
- Default no-train
- Word-integrated workflow lowers adoption friction

WEAKNESSES

- Less mature trust portal
- No explicit US-only residency configuration
- Subprocessor list NDA-gated
- ABA Op 512 mapping less detailed than top-tier legal vendors

BEST USE CASE

Small-to-mid firms (5-50 attorneys) focused on transactional / contract work, where Word-integration and per-attorney pricing match the budget and workflow.

AVOID WHEN

Firms with strict regulatory scrutiny (especially BigLaw or in-house teams under heavy compliance scrutiny) that need top-tier trust documentation.

Last reviewed: 2026-05-13 · Homepage: https://www.spellbook.legal

12-AXIS SCORING

Axis	Status	Note	Source
BAA / DPA available	Yes	Heidi signs BAAs for US enterprise customers.	Heidi Security
Training-data opt-out	Yes	Heidi does not train models on customer encounter data.	Heidi Privacy
US data residency option	Partial	Heidi offers US-region hosting for US customers. Default configuration may use multi-region infrastructure; explicit US-only residency requires enterprise contract.	Heidi Security
SOC 2 Type II report	Partial	Heidi reports SOC 2 audit completion; report distribution via direct enterprise request.	Heidi Security
ISO/IEC 42001 attestation	No	No ISO/IEC 42001 attestation as of May 2026.	Public posture review
NIST AI RMF self-attestation	No	No public NIST AI RMF self-attestation. Heidi's primary regulatory anchoring is Australian (TGA) given its origin market.	Public posture review
Colorado AI Act readiness	No	No Colorado AI Act-specific public statement.	Public posture review
HHS-OCR Section 1557 readiness	Partial	Heidi documents general clinical safety; explicit Section 1557 public statement less developed than US-headquartered peers.	Heidi documentation
FRB SR 11-7 readiness	N/A	Healthcare-vertical positioning.	Heidi positioning
ABA Formal Op 512 readiness	N/A	Healthcare-vertical positioning.	Heidi positioning
Subprocessor list public	Partial	Subprocessor information available on request; not self-serve public.	Heidi Security
Trust-center maturity	2/5	Security documentation present but less mature than US-headquartered peers. AI-specific governance for US market expanding but behind Abridge / Suki / DAX.	heidihealth.com/security

DEEP DIVE

Heidi is the price-leader in clinical AI documentation — meaningfully cheaper than DAX Copilot, Abridge, or Suki at small-practice scale. The governance posture reflects the smaller-vendor scale and the Australian origin: BAA available but trust-portal maturity and US-regulatory-specific documentation (Section 1557, Colorado AI Act, NIST AI RMF) are less developed than US-headquartered peers.

STRENGTHS

- BAA-eligible
- Significantly lower price point than US-headquartered peers
- Default no-train

WEAKNESSES

- Trust portal less mature than US peers
- Section 1557 documentation less developed
- No NIST AI RMF or Colorado AI Act statement
- Explicit US-only residency requires enterprise contract

BEST USE CASE

Solo and small practices (1-15 providers) where price sensitivity is high and the governance burden is correspondingly smaller (lower OCR scrutiny than a multi-state health system).

AVOID WHEN

Health systems, hospital networks, or any organization under active OCR Section 1557 scrutiny. The trust-portal maturity gap and weaker public US-regulatory engagement create defensibility risk during audit.

Last reviewed: 2026-05-13 · Homepage: https://www.heidihealth.com

12-AXIS SCORING

Axis	Status	Note	Source
BAA / DPA available	No	Notion does not sign BAAs. Notion has explicitly stated it is not HIPAA-compliant and should not store PHI.	Notion HIPAA support article
Training-data opt-out	Partial	Notion AI does not train on workspace content by default for Business and Enterprise plans. Free and Plus: opt-out toggle available.	Notion AI Privacy
US data residency option	No	No US data residency configuration option as of May 2026. Notion uses AWS US-East default.	Notion Trust Center
SOC 2 Type II report	Yes	SOC 2 Type II report available via Notion Trust Center under NDA. ISO 27001:2022 also held.	Notion Trust
ISO/IEC 42001 attestation	No	No ISO/IEC 42001 attestation.	Public posture review
NIST AI RMF self-attestation	No	No public NIST AI RMF self-attestation.	Public posture review
Colorado AI Act readiness	No	No Colorado AI Act compliance statement.	Public posture review
HHS-OCR Section 1557 readiness	N/A	Not BAA-eligible — Section 1557 use case disqualified by HIPAA gap.	HHS-OCR Section 1557 — deployer scope
FRB SR 11-7 readiness	N/A	SR 11-7 is deployer responsibility for banking use, but the lack of BAA already disqualifies most regulated bank deployments.	FRB SR 11-7 — deployer scope
ABA Formal Op 512 readiness	N/A	ABA Op 512 is practitioner responsibility; no BAA significantly raises the privilege bar for law firm use.	ABA Formal Op 512 — practitioner scope
Subprocessor list public	Yes	Notion subprocessor list public (OpenAI as Notion AI subprocessor, AWS, Stripe, etc.).	Notion Subprocessors
Trust-center maturity	3/5	Mature trust portal with SOC 2 + ISO under NDA. AI-specific governance documentation is thin — no Colorado AI Act, no NIST AI RMF, no ISO 42001.	Notion Trust

DEEP DIVE

Notion AI is one of the most-deployed shadow-AI vectors in the regulated mid-market. The product is good and widely loved — but the lack of BAA, lack of residency, and thin AI-specific governance documentation make it a poor fit for any regulated workload. Most firms we audit have Notion AI in use and PHI/PII in Notion without realizing the BAA gap.

STRENGTHS

- No-train default for Business/Enterprise
- Mature SOC 2 + ISO 27001 posture
- Public subprocessor list

WEAKNESSES

- No BAA — not HIPAA-compliant
- No US data residency option
- No AI-specific governance documentation
- Common shadow-AI vector for regulated data

BEST USE CASE

Non-regulated workspace use where no PHI, PII, or privileged data enters Notion. Internal-only knowledge management for non-regulated workloads.

AVOID WHEN

Any environment where PHI, regulated financial data, or privileged legal content might enter a Notion workspace. DLP at the email/upload boundary is the right preventive control.

Last reviewed: 2026-05-13 · Homepage: https://www.notion.so/product/ai · Trust center: https://www.notion.so/help/notion-trust

12-AXIS SCORING

Axis	Status	Note	Source
BAA / DPA available	No	Meta does not offer a BAA directly. BAA must be obtained from the hosting partner (AWS Bedrock, Azure AI Studio, GCP Vertex) where Llama is deployed. Self-hosted deployments shift the entire BAA burden to the deploying organization.	Meta Llama Community License
Training-data opt-out	Yes	Open weights — no training feedback loop to Meta. Inputs to your hosted deployment never leave your tenant.	Meta Llama license terms
US data residency option	Yes	Self-hosted or partner-hosted on a US region — deploying organization controls residency entirely.	Deployment-controlled
SOC 2 Type II report	No	Meta does not provide SOC 2 for Llama directly. Hosting partner (AWS/Azure/GCP) provides cloud-side SOC 2.	Meta Trust Center
ISO/IEC 42001 attestation	No	No ISO/IEC 42001 attestation.	Public posture review
NIST AI RMF self-attestation	No	No NIST AI RMF self-attestation. Meta publishes Responsible Use Guide and Model Card; deploying organization performs RMF mapping.	Meta Responsible Use Guide
Colorado AI Act readiness	No	No Colorado AI Act compliance statement. Deployer responsibility entirely.	Public posture review
HHS-OCR Section 1557 readiness	N/A	Foundation model — Section 1557 is deployer responsibility.	HHS-OCR Section 1557 — deployer scope
FRB SR 11-7 readiness	N/A	Foundation model — SR 11-7 is deployer responsibility.	FRB SR 11-7 — deployer scope
ABA Formal Op 512 readiness	N/A	Foundation model — ABA Op 512 is deployer responsibility.	ABA Formal Op 512 — practitioner scope
Subprocessor list public	No	Self-hosted: no Meta subprocessor chain. Partner-hosted: hosting partner's subprocessor list applies.	Deployment-controlled
Trust-center maturity	2/5	Meta publishes Responsible Use Guide, model cards, license terms. No trust portal in the OpenAI/Anthropic sense. Compliance posture lives at the hosting layer.	llama.com

DEEP DIVE

Llama scores poorly on a vendor-governance scorecard because Meta delegates governance to the deploying organization. This is by design — open weights mean the deployer owns the entire stack. The right way to evaluate Llama is to score the hosting partner (AWS Bedrock, Azure AI, Vertex AI) instead, because that's where the BAA, SOC 2, residency, and subprocessor controls actually live.

STRENGTHS

- Open weights — full deployer control of data, residency, retention
- No training feedback loop to Meta
- Cost advantage at scale via self-hosting

WEAKNESSES

- No vendor-side BAA, SOC 2, residency, or subprocessor controls
- Deployer owns 100% of governance burden
- No NIST AI RMF self-attestation, no Colorado AI Act statement

BEST USE CASE

Organizations with mature ML/AI platform teams that need full data control, are running on-prem or sovereign-cloud workloads, or have validated hosting on AWS Bedrock / Azure AI Studio / GCP Vertex with the hosting partner's BAA in place.

AVOID WHEN

Smaller organizations without an internal AI platform team. The cost of building deployer-side governance on top of Llama exceeds the cost of paying for OpenAI Enterprise or Claude for Work in most mid-market scenarios.

Last reviewed: 2026-05-13 · Homepage: https://llama.com

12-AXIS SCORING

Axis	Status	Note	Source
BAA / DPA available	No	Otter.ai does not currently offer a BAA. Otter has stated HIPAA compliance is not supported.	Otter.ai Security FAQ
Training-data opt-out	Partial	Enterprise tier: customer audio/transcripts not used for model training. Free/Pro: opt-out toggle available; defaults vary by feature.	Otter Privacy Policy
US data residency option	No	No documented US data residency configuration as of May 2026.	Public posture review
SOC 2 Type II report	Yes	SOC 2 Type II completed; report available via direct request.	Otter Security
ISO/IEC 42001 attestation	No	No ISO/IEC 42001 attestation.	Public posture review
NIST AI RMF self-attestation	No	No public NIST AI RMF self-attestation.	Public posture review
Colorado AI Act readiness	No	No Colorado AI Act compliance statement.	Public posture review
HHS-OCR Section 1557 readiness	N/A	Not BAA-eligible — disqualifies clinical use.	HHS-OCR Section 1557 — deployer scope
FRB SR 11-7 readiness	N/A	SR 11-7 is deployer responsibility.	FRB SR 11-7 — deployer scope
ABA Formal Op 512 readiness	N/A	Practitioner responsibility; lack of BAA significantly raises privilege risk for law firm use.	ABA Formal Op 512 — practitioner scope
Subprocessor list public	Partial	Subprocessor list available to enterprise customers on request. Not self-serve public.	Otter Security FAQ
Trust-center maturity	2/5	Security page exists but is thin. AI-specific governance documentation absent. Lower-maturity trust posture.	otter.ai/security

DEEP DIVE

Otter.ai is widely deployed in sales/CS organizations and routinely creeps into clinical, financial, and legal meeting workflows without governance review. The product is competent; the governance posture is not aligned to regulated use. The most common audit finding involving Otter is patient or attorney-client conversations transcribed without a BAA or privilege protocol.

STRENGTHS

- SOC 2 Type II
- Enterprise no-train default
- Mature transcription product

WEAKNESSES

- No BAA, no HIPAA support
- No US residency option
- Thin AI-specific governance documentation
- Subprocessor list not self-serve public

BEST USE CASE

Non-regulated meeting transcription — sales call notes, internal team meetings, marketing planning sessions.

AVOID WHEN

Patient encounters, attorney-client conversations, confidential financial advisory meetings. Use a BAA-covered alternative (Microsoft Teams transcription under M365 BAA, or sector-specific tools like DAX Copilot).

Last reviewed: 2026-05-13 · Homepage: https://otter.ai · Trust center: https://otter.ai/security

12-AXIS SCORING

Axis	Status	Note	Source
BAA / DPA available	No	No BAA available as of May 2026 — Perplexity is not a HIPAA business associate. Do not use for PHI workflows.	Perplexity Privacy Policy
Training-data opt-out	Partial	Enterprise Pro contract terms exclude customer data from training. Consumer tiers: opt-out available via account settings.	Perplexity Enterprise Privacy
US data residency option	No	No documented US data residency configuration for enterprise customers as of May 2026.	Public posture review
SOC 2 Type II report	Partial	Perplexity has publicly claimed SOC 2 Type II completion. Report distribution via direct request, not a self-serve trust portal.	Perplexity Enterprise security page
ISO/IEC 42001 attestation	No	No ISO/IEC 42001 attestation.	Public posture review
NIST AI RMF self-attestation	No	No public NIST AI RMF self-attestation.	Public posture review
Colorado AI Act readiness	No	No Colorado AI Act compliance statement.	Public posture review
HHS-OCR Section 1557 readiness	N/A	Section 1557 is deployer responsibility for any clinical use — but the absence of a BAA makes Perplexity unsuitable for PHI use cases.	HHS-OCR Section 1557 — deployer scope
FRB SR 11-7 readiness	N/A	SR 11-7 is deployer responsibility for any banking use.	FRB SR 11-7 — deployer scope
ABA Formal Op 512 readiness	N/A	ABA Op 512 is practitioner responsibility for any legal research use.	ABA Formal Op 512 — practitioner scope
Subprocessor list public	Partial	Perplexity uses multiple model vendors as subprocessors (OpenAI, Anthropic, Mistral). Subprocessor list available to enterprise customers under NDA.	Perplexity Enterprise Privacy
Trust-center maturity	2/5	No self-serve trust portal. Enterprise security documentation available on request. Material gap for regulated buyers.	Perplexity Enterprise

DEEP DIVE

Perplexity is best understood as an answer-engine layer that fans out to multiple foundation models behind the scenes. The governance gap is structural: Perplexity inherits some posture from upstream models but doesn't sign HIPAA BAAs and doesn't publish a Colorado AI Act / NIST AI RMF posture. Strong for general research, weak for regulated workflows.

STRENGTHS

- Citation-grounded responses reduce hallucination risk vs. raw chat
- Enterprise contract excludes customer data from training
- SOC 2 Type II claim

WEAKNESSES

- No BAA — disqualifies for PHI
- No US data residency option
- No NIST AI RMF, ISO 42001, or Colorado AI Act statement
- No self-serve trust portal

BEST USE CASE

General-purpose research use cases where the citation-grounded format is a real advantage and no regulated data is involved.

AVOID WHEN

Any PHI, regulated financial data, or privileged legal content. Do not deploy in clinical, banking, or law firm production workflows without an alternative.

Last reviewed: 2026-05-13 · Homepage: https://www.perplexity.ai