




Safe & Trusted AI For Enterprises

Aligned with Governance,
Explainability, and Scale





AI Systems Enterprises Can Trust

This is where innovation meets discipline. Cygnet.One helps enterprises build and deploy AI systems that perform in production without compromising governance, privacy, or compliance.

Our AI solutions are engineered for real-world environments where decisions must be explainable, auditable, and aligned with enterprise risk frameworks.

◆ AI Risk and Compliance Framework ◆

Governed AI

Controls, approval gates, and audit trails built in

Production-Ready AI

Monitoring, drift detection, and operational reliability

Secure by Design

Privacy, access controls, and traceability across workflows

Compliance-Ready Outputs

Explainable decisions that stand up to review

Our Methodology



AI adoption is accelerating faster than governance can keep pace. Enterprises want speed and intelligence, but they cannot afford opacity, uncontrolled autonomy, or compliance risk.

At **CYGNET.ONE**, Safe & Trusted AI is not a statement. It is an engineering discipline.

Our approach is anchored in the **SAILER** Framework (**S**afe **A**I which is **L**awful **E**thical **R**esponsible), a lifecycle methodology that governs how AI is designed, built, deployed, and continuously monitored in real enterprise environments.



The Foundation Layer: Lawful AI

- Compliance with laws, regulations, and standards
- Privacy, data protection, and sectoral requirements
- Consent, purpose limitation, and retention alignment
- Cross-border and residency compliance

The Value Layer: Ethical AI

- Fairness and non-discrimination
- Avoidance of harmful outcomes and unintended bias
- Respect for human autonomy and dignity
- Alignment with organisational values and societal norms

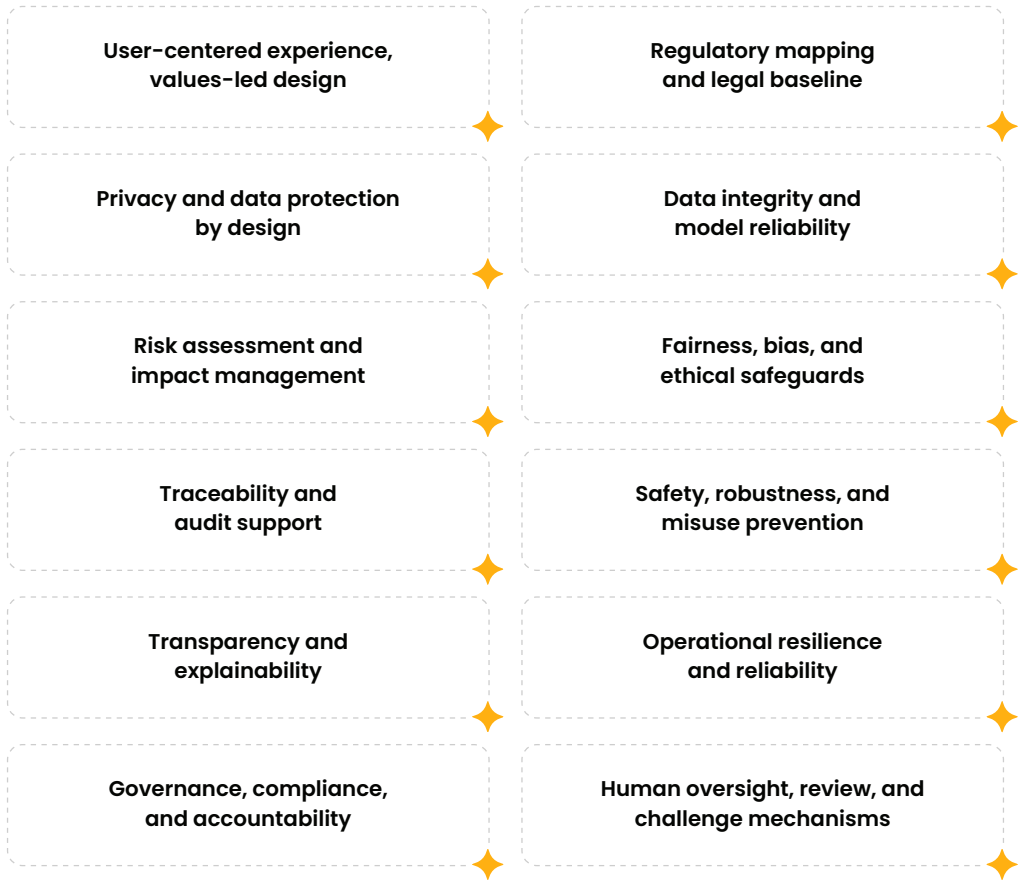
The Control Layer: Responsible AI

- Governance frameworks and decision accountability
- Risk identification, monitoring, and mitigation
- Model lifecycle controls (train, test, deploy, monitor)
- Auditability, incident management, and human-in-the-loop safeguards

Principles of Safe & Trustworthy AI



The SAILER Framework Pillars



The CYGNET.ONE Methodology



Explore Technical Assets

www.cygnnet.one | hello@cygnnet.one


Scan Now



Why This Matters

Enterprise AI must do more than perform well. It must operate within clear boundaries, explainable logic, strong data controls, and governance that stands up to compliance and audit review.

Agents as a Service

Controlled AI Agents for Enterprise Workflows

Task-driven agents built with defined scope, permission boundaries, and clear accountability.

Guardrails include

- Controlled data access and approved execution paths
- Restricted decision authority
- Continuous monitoring of agent behaviour

Typical use cases

Operational automation, decision support, and predictable enterprise workflows.

HireAI

AI Interviewer with Built-In Fairness and Governance

An AI-powered interviewer for structured hiring, consistent evaluation, and compliance-aligned interview processes.

Risk controls include

- Role-specific evaluation criteria
- Standardized scoring and explainable evaluation logic
- Controlled prompts to reduce inconsistency
- Human oversight remains central

Business outcome

Faster interview cycles, higher consistency, and ethical hiring support at scale.

Conversational AI

AI-Powered Voice Assistant You Can Trust

Real-time voice interactions for customer service, support, and enterprise conversational workflows.

Safety and compliance controls include

- Secure handling of voice data and encryption for sensitive interactions
- Policy-driven response behaviour
- Escalation rules for high-risk or uncertain queries

Business outcome

Accurate conversations, safer responses, and compliance-ready voice experiences

Common Controls Across All Solutions

- Clear boundaries on what AI can and cannot do
- Explainability for critical decisions
- Continuous monitoring of behaviour and outputs
- Secure data handling across interactions

Designed for Enterprise Reality

Built for regulated environments, audit readiness, modern cloud architectures, and scale without compromising control.

Explore Technical Assets

www.cygnet.one | hello@cygnet.one


Scan Now



AI Platforms Ready for Regulated & High-Risk Environments



Where accuracy and auditability are non-negotiable

Finance, legal, and tax operations demand precision. Cygnet.One builds AI systems that support decisions without bypassing responsibility.

Enterprises manage high volumes of documents, complex decisioning, and strict regulatory mandates. Traditional automation struggles to scale, adapt, and remain auditable. Our AI solutions combine advanced intelligence with governance so every output is **explainable, traceable, and compliance-aligned**.

Intelligent Document Processing (VLM-Powered)

Documents understood, not just digitized

OCR alone cannot reliably handle contracts, invoices, tax forms, and financial statements. Our Vision-Language Model powered IDP interprets layout, context, semantics, and business intent.

Core capabilities

- Context-aware extraction with validation rules
- Confidence scoring to flag uncertain results
- Integration with ERP, tax platforms, and data lakes

Trusted by design

- Compliance logic embedded into extraction
- Human review triggered by confidence thresholds
- Full traceability from source document to structured output

Business impact

- Up to **75% reduction** in manual document review time
- Fewer exceptions and rework
- Higher accuracy with compliance checkpoints

Intelligent Accounts Receivable/Accounts Payable

Predictive finance with control

Finance teams often lack foresight, leading to late collections and working capital inefficiencies. Our AI-driven AR/AP solution combines document intelligence with prediction models for explainable cashflow decision support.

What it does

- Predicts receivable inflows using historical patterns and risk profiles
- Optimizes payables using vendor data, compliance status, and policy rules
- Provides explainable recommendations with dashboards and approval trails

Business outcomes

- Faster working capital cycles
- Reduced DSO
- More confident forecasting for CFOs and boards

AI-First Litigation & GST Intelligence

Legal and tax intelligence with accountability

Manual litigation and compliance processes increase the risk of missed deadlines, inconsistent responses, and unmanaged exposure. Our platform centralizes notices, cases, deadlines, obligations, and outcomes.

Key features

- AI-based intake, classification, and extraction of key elements
- Risk and impact scoring using rule libraries and historical data
- AI suggestions for responses and next steps
- Maker-checker controls and leadership dashboards

Designed for compliance

- Source references attached to AI suggestions
- Human legal teams retain final approval
- Full audit trails for regulators and internal governance

Enterprise benefits

- Reduced backlog
- More consistent response quality
- Early visibility into exposure and provisioning needs

Explore Technical Assets

www.cygnnet.one | hello@cygnnet.one


Scan Now



Technology Services for AI-Led Enterprise Growth



Where accuracy and auditability are non-negotiable

Cygnet.One helps organizations simplify complex systems into technology that delivers measurable business outcomes. We combine expertise in **data, quality engineering, cloud, and managed services**, with an AI-first mindset so teams can deliver faster and with confidence.

Data and AI Services

Reliable data foundations and AI-ready intelligence.

Core capabilities

- Data engineering and management (pipelines, architectures, integration)
- Data migration and modernisation (safe legacy-to-modern moves)
- Business analytics and embedded AI (dashboards, predictive insights, decision support)
- Insights-driven transformation aligned to business outcomes

Business impact

- Unified data foundation, predictive insights, AI-ready platforms.

Quality Engineering

Quality built into the lifecycle to reduce risk and speed delivery.

Core capabilities

- Test consulting and maturity assessment
- Business and functional testing
- End-to-end enterprise testing
- Data transformation testing

Business impact

- Higher release quality, lower defect leakage, stronger delivery confidence.

Managed IT Services

Stable, secure operations aligned to enterprise goals.

Core capabilities

- IT strategy and consulting
- Application and infrastructure support
- Cybersecurity and compliance
- Governance, risk management, and compliance (GRC)

Business impact

- Reduced downtime, controlled risk, teams focused on priorities.

Cloud and DevOps Services

Cloud environments designed to scale, optimize, and support AI.

Core capabilities

- Cloud strategy and architecture
- Migration, modernisation and Cloud-native development
- AWS Advanced Tier Partner
- Operations, monitoring, and optimization

Business impact

- Resilient platforms, reduced overhead, AI-ready infrastructure.

AI Solutions and Platforms

Enterprise-grade AI products built with governance, audit trails, and human oversight options. These are engineered systems designed to operate safely within enterprise controls.

Tax Datalake

Risk Engine

Conversational AI

Agents as a Service

Apps in ChatGPT

HireAI

Explore Technical Assets

www.cygnnet.one | hello@cygnnet.one


Scan Now




Case Studies:

Real Results. Real Impact. Proven Outcomes



HireAI – AI-First Talent Intelligence

Enterprises hiring at scale struggle with resume overload, inconsistent shortlisting, and role-fit subjectivity, especially for niche tech and domain roles.



How HireAI Helps


- AI-driven resume parsing & skill mapping
- Role-fit scoring with bias-aware evaluation
- Faster, consistent shortlisting at scale

Impact

- 50–60% faster hiring cycles
- 40% reduction in recruiter effort

AI-First Accounts Receivable & Accounts Payable Intelligent Bot

Finance teams spend disproportionate time pulling data from ERP, reconciling mismatches, and answering ad-hoc cash-flow questions—not analyzing outcomes.



How the AR & AP Bot Helps

- Conversational access to AR/AP insights
- Real-time collections, aging & risk visibility
- Automated reconciliation & exception detection

Impact

- 15–25% improvement in DSO
- 60–70% reduction in manual finance analysis

AI-First Litigation Management System (LMS)

Tax and legal teams manage hundreds of notices across jurisdictions, with manual tracking, scattered documents, and deadline anxiety driving compliance risk.



How LMS Helps


- AI-based notice parsing & issue extraction
- Contextual classification using laws & precedents
- Centralized case repository with alerts & risk views

Impact

- 75% faster notice handling
- 50% reduction in manual legal effort

Agentic AI Use Case – Cygnet Pre-Sales Assistant

Sales teams lose momentum due to slow lead follow-ups, repeated product queries, and manual pre-sales effort, especially in high-volume inbound scenarios.



How the Pre-Sales Agent Helps

- Conversational lead engagement & qualification
- Product/service discovery & intent detection
- Proposal and quotation assistance

Impact

- 40% faster lead response time
- Improved lead-to-opportunity conversion

Explore Technical Assets

www.cygnet.one | hello@cygnet.one


Scan Now



AI-Powered Risk Engine: Predict, Prevent, and Protect in Real Time



Why This Matters

AI is now driving high-impact decisions across BFSI, insurance, healthcare, and digital platforms. Once AI goes live, unmanaged risk becomes immediate business exposure.

What the AI Risk Engine Does

A continuous monitoring and governance layer that prevents AI risk from becoming a production incident. It makes AI decisions **observable, traceable, and reviewable** across data, models, and operations.

Core Risk Areas Covered

Model Risk

- Accuracy decay and drift detection
- Abnormal behaviour alerts in production



Data Risk

- Data lineage and usage visibility
- Quality checks across training and inference pipelines

Bias and Fairness Risk

- Bias detection across protected attributes
- Statistical + rule-based checks
- Explainability support for governance reviews



Operational and Compliance Risk

- AI usage mapped to internal policies
- Regulatory alignment monitoring
- Audit-ready evidence trails

Built for AWS Environments

■ Designed for AWS-native AI and data services, with integration support for:

Amazon SageMaker

AWS Glue

Amazon QuickSight

AWS Lake Formation

Amazon CloudWatch

AWS IAM

Monitoring and Security (AWS-Ready)

- Event-driven monitoring for real-time detection
- Secure logging and traceability
- Encryption for sensitive model and data artifacts
- Designed to minimize impact on inference workloads

Who It's Designed For

Enterprises running AI workloads on AWS, especially regulated industries where AI governance, audit readiness, and compliance cannot be optional.

What You Gain

- Continuous AI risk visibility
- Explainable and auditable decisions
- Stronger regulatory alignment
- Trustworthy AI systems at scale

Explore Technical Assets

www.cygnnet.one | hello@cygnnet.one


Scan Now



About us

CYGNET.ONE is an AI-first technology services provider with 25 years of enterprise delivery experience. We help organizations design, build, modernize, and operate digital systems across data, cloud, quality engineering, managed services, and governed AI solutions.

Our Safe & Trusted AI approach is anchored in the **SAILER Framework**, supported by enterprise-grade controls such as traceability, human oversight, and continuous monitoring.

CYGNET.ONE supports enterprises across regulated industries including BFSI, insurance, healthcare, and large digital platforms, delivering outcomes from strategy through execution.

Global Presence:

Across 35+ Countries

Global Markets:

Americas | UK & Europe | Africa | Middle East | Asia Pacific

Email us at

hello@cygnet.one

Discover more

www.cygnet.one



Scan to know more