DATA SHEET

C3.ai

C3 Agentic AI Platform

Enterprise AI Platform for Rapidly Developing, Deploying, and Operating Enterprise-Scale AI Applications

The C3 Agentic AI Platform is a software platform that uses a flexible, model-driven architecture to accelerate and reduce the complexities of developing enterprise-grade AI applications. With C3 Agentic AI Platform, organizations can integrate advanced AI capabilities seamlessly and deliver AI-enabled applications faster than alternative methods to capture value quickly.

C3 AI Named a Leader in The Forrester Wave™: AI/ML Platforms, Q3 2024

FORRESTER®
WAVE LEADER 2024
AI/ML Platforms

C3 AI scores #1 in:

Vision
Innovation
Model Development
Model Governance
Data Governance
Security
Partner Ecosystem

"C3 AI offers an AI future to enterprises...C3 AI is a good fit for enterprises that want a solid AI platform for bespoke AI applications that also includes a bountiful library of pre-built applications."

C3 AI Studio

C3 AI Studio offers a rich library of deep-code tools and a low-code environment for developing, deploying, and operating enterprise AI applications. C3 AI Studio provides a cohesive development experience on a visual canvas providing data ingestion, data modeling, machine learning feature engineering and model lifecycle management, and metadata-driven UI development tooling. C3 AI Studio allows developers and data scientists to focus on solving complex business problems by providing an integrated environment that abstracts routine and complex application development tasks.

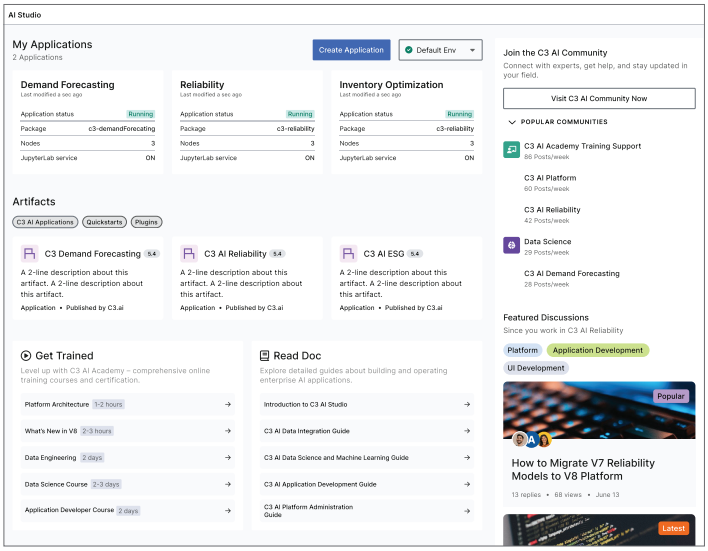

Agentic AI Services and Capabilities

The C3 Agentic AI Platform offers advanced agentic AI services and capabilities, including omni-modal data fusion, secure deployment, advanced orchestration, and a robust agent and tools framework. These features enable seamless integration and management of multiple AI models within a unified workflow, supporting a wide range of high-value, cross-industry applications and enhancing operational efficiency, security, and decision-making across sectors.

Seamless AI Model Fusion

Leverage advanced omni-model capabilities to deliver highly accurate results. Seamlessly integrate and orchestrate multiple AI models within a unified workflow, enabling support for a broader spectrum of tasks and delivering comprehensive solutions. Execute sophisticated retrieval and reasoning operations to address complex problems and inform decision-making across diverse, high-value, cross-industry applications.

Secure Cloud AI Deployment

Deploy AI models within a secure cloud environment, whether public or private, ensuring that all data and operations are contained within a protected enclave. Enable comprehensive management and oversight of the AI models' operations, including data handling, processing, and storage, while strictly adhering to security and privacy policies.

Flexible Agent and Tool Framework

Create intelligent, autonomous agents capable of performing a wide range of tasks with minimal human intervention. Support the integration of various tools and APIs, enabling agents to interact with external systems, retrieve information, and execute tasks. Agents can utilize application data models and APIs and custom retrieved context to provide relevant and accurate responses, enhancing their ability to perform complex tasks and support decision-making.

Context Aware, AI-Driven Applications

Develop context-aware applications that understand and respond to the context in which they operate. Combine data from various sources, such as databases, sensors, or user inputs to automate tasks, make predictions, and provide essential insights to application users.

Multi-Agent Workflows

Design and implement workflows that involve one or more AI agents working together. Agents can perform different tasks or collaborate on a common goal, enhancing the flexibility and complexity of the workflows. Integrate human oversight into the application workflow. Pause processes to allow for a human review, validation, or decision-making before proceeding. Ensure human expertise and judgment are incorporated where necessary to improve the overall reliability and accuracy of the system.

Observability

Monitor the inner workings of agentic workflows with out-of-the box tracing for LLM requests, tools, and pipelines across multiple services. Track latency, throughput, resource utilization, and other key metrics of individual steps to identify performance bottlenecks and quickly diagnose issues.

Visual Studio Code Extension

C3 AI Studio provides code-based experiences for technical users through an extension of Visual Studio Code source-code editor. Developers can use C3 AI-specific IntelliSense, leveraging out-of-the-box auto-suggestions and autocompletion across C3 AI models. Data scientists can write custom Python methods and inspect any issues in their logic using an integrated Python debugger and QA engineers can manage test files across multiple applications.

Develop applications in a familiar IDE setting

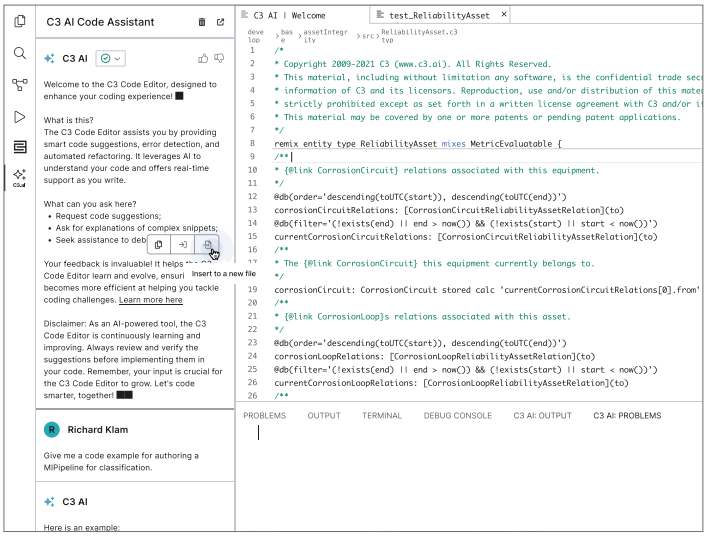

- Leverage and integrated AI-powered coding assistant that provides code generation and documentation assistance.
- Instantly check for errors on data model construction and function implementation with each file save
- Hover over any keyword, data model element or function to receive in-context documentation
- Click into references and implementation files with the code to traverse across an application and its dependencies

JupyterLab

C3 AI's on-demand JupyterLab Notebooks provides an interactive interface to all data and machine learning services to enable a complete data science lifecycle.

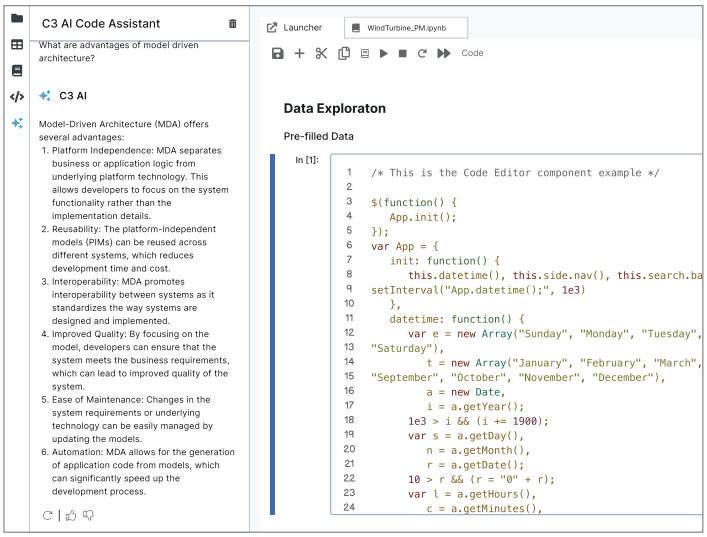

- Integrate generative AI
- Train and deploy models developed in Jupyter or distribute intensive workloads to an auto-scaling compute cluster
- Take advantage of dedicated compute environments backed by configurable CPU and GPU profiles
- Share notebooks across all data scientists in the application, and optionally commit notebooks to the code repository for CI testing and reuse across applications

Application Canvas

C3 AI Studio exposes the power of C3 AI's model-driven architecture through an intuitive application canvas and provides an integrated set of best-of-breed tools that allow business experts, data scientists, data engineers, application developers and IT professionals to easily collaborate on developing, deploying, and operating complex AI applications.

Data

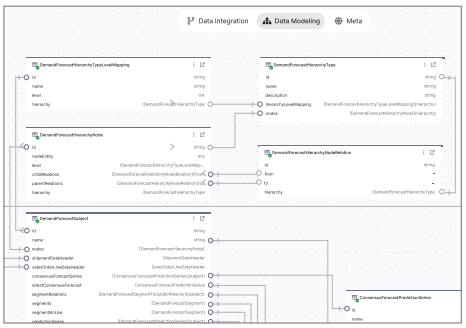

- Seamlessly integrate disparate internal and external data
- Explore, clean, contextualize, and label data through human-in-the-loop and AI-driven approaches
- Visualize data pipelines and develop features through an integrated code experience, drastically reducing the effort required to move data pipelines into production

Machine Learning

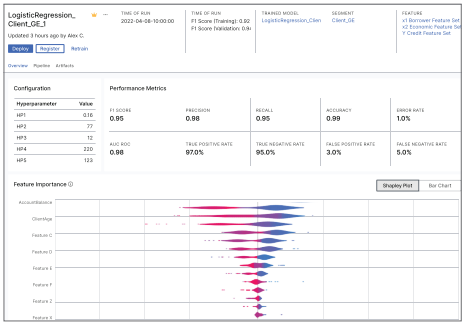

- Access all upstream and downstream workflows like data engineering and application user interface development
- Configure, run, and track ML experiments
- Connect to JupyterLab from the application canvas with one-click

Composable UI

- Leverage low-code tools to easily configure applications to business-specific needs
- Theme applications to make sure AI applications match branding guidelines
- Develop new processes and workflows by leveraging low-code tools
- Use modern development frameworks such as React to dynamic and responsive user interfaces

DevSecOps

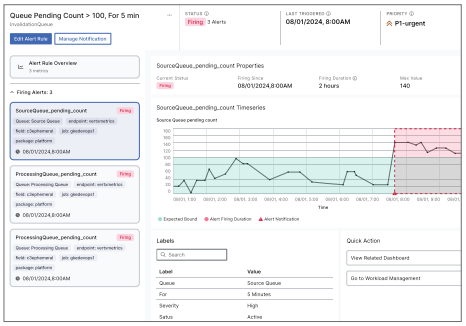

- Configure continuous integration (CI) pipelines, inspect the qualities of build artifacts, and deploy green builds into production
- Monitor the health of environments, configure environment settings, and manage usage and resourcing
- Review application build quality and dive deep to diagnose the causes
- Promote release candidates into production safely and analyze the health of deployments
- Leverage a shared repository of C3 AI Applications, using specific data modules, ML modules, or UI modules to help easily extend applications

AI/ML Platform

C3 Agentic AI Platform provides a collaborative development environment for data science teams to rapidly prototype AI / ML models, manage experiments, deploy models into production, and manage model lifecycles at scale. Users can run pre-packaged experiments to automate time-intensive tasks such as parallelized training and hyperparameter optimization and scale up model training using configurable hardware and an auto-scaling compute cluster. C3 Agentic AI Platform automatically records experiments and training progress and provides an interactive interface for comparing the performance scores of all models, simplifying the discovery of leading models.

Model Governance

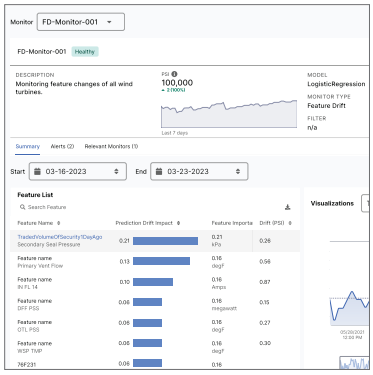

- Native model inference service for secure, audited serving of C3 AI and open-source models, including large language models
- C3 AI Model Registry provides a central repository to govern the end-to-end ML Lifecycle
- Configure approval and review workflows to ensure models meet the organizational and regulatory requirements
- Manage the entire model governance lifecycle, from deployment to monitoring and improvement
- Automatically track hyperparameters, metrics, algorithms, and data sources to identify the best model
- Deploy models through code or no-code tools
- Continuously test models and monitor model performance

Model Development

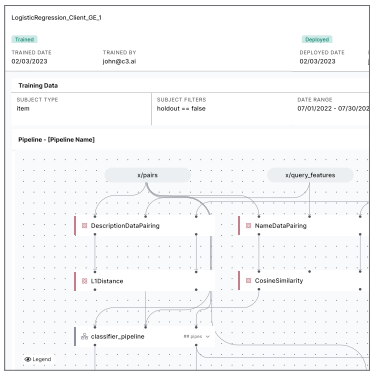

- Compose multi-step algorithms using pre-packaged or custom-defined C3 AI ML Pipes with automated scoring and explainability to accelerate the development and operationalization of complex AI use cases
- Support for open source and commercial LLMs including Azure Open AI, AWS Bedrock, Google Gemini, Anthropic Claude, Meta Llama 3, Mistral, and others
- Use 15 pre-built explainers and 30 pre-built scoring metrics or customize the techniques for each ML Pipeline.
- Utilize popular ML libraries (e.g., Transformers, Tensorflow, XGBoost) or integrate custom libraries and techniques
- LLM agents and tools for generative AI application development: orchestration, retrieve-extract-answer pipelines, database queries, geospatial/math, planning/execution

Data Governance

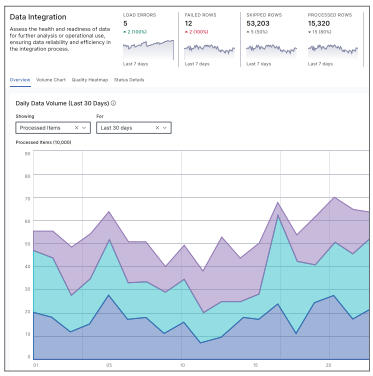

- Integrate any source - timeseries, relational, image, video - via batch/stream, or virtualization using connectors from library of 250+ external source connectors
- Visualize the application data model and lineage
- Visually profile petabyte-scale data using over 120 analytic operators, 50+ statistical operators, and 20+ configurable visualizations
- Define ML features using pandas or an optimized time series function library from C3.ai
- Develop and test features interactively in Jupyter with integrated metadata that provides advanced lineage and associates all source data, feature computations, and snapshots used to train a model
- Seamlessly deploy to features to production for ongoing, parallelized, feature materialization

Additional Platform Services and Capabilities

The C3 Agentic AI Platform provides a comprehensive set of services and capabilities to help organizations build and deploy AI applications faster than alternative methods. The platform's model-driven architecture, data integration, management and processing capabilities, time series services, AI and model management, and security framework accelerate data science and application development to help enterprises capture value quickly.

Model-Driven Architecture

Enable greater data science and application developer productivity, rapidly deliver enterprise-scale AI applications, and future-proof existing IT investments. The C3 Agentic AI Platform uses conceptual models of all the attributes and processes related to a specific entity or domain as well as physical objects or data stores. C3 AI's model-driven architecture can represent application data, metadata, processes, interrelationships, persistence, computing processes, time series expressions, language bindings, and AI/ML tools and algorithms.

Data Integration Services

Enable rapid integration of data from enterprise, extraprise, and sensor data feeds with support for both structured and unstructured data. The C3 Agentic AI Platform can ingest data in batch, stream, or message-based integrations. C3 AI has prebuilt connectors to many common data sources including Postgres, Oracle, SAP, HBase, HDFS, Apache Kafka, AWS Kinesis, OSI PI, and Cassandra. Data integration services are extensible, enabling developers to configure and enable additional connectors.

Data Management Services

Enable persistence of large volumes of data, while also making data readily available for analytical calculations. Virtualize external data stores within the C3 Agentic AI Platform for AI algorithms and applications. Data management services include data federation, management of and interaction with multiple databases, and persistence of data in the appropriate data store.

Time Series Services

Enable persistence, processing, and representation of data objects as time series, including the ability to normalize or calendarize data (e.g., time-align data, retrieve time series at different time intervals), identify and flag gaps in data, manage data that are received out of sequence, and apply pre-built (or custom) mathematical expressions on time series data. Seamlessly manage time series data and costs across hot (fast reads/writes, higher costs) and cold (slower reads/writes, lower costs) storage.

AI and Model Management Services

Manage models across machine learning life cycle stages, including model design and experimentation, model training and evaluation, model integration and deployment, production inference, and model maintenance.

Security

Deliver end-to-end authentication and authorization, including access control to data and methods, using the role-based, and certified C3 AI security framework. SOC2, SOC3, NIST, and HIPAA attestations.

Multi-cloud and Edge Deployments

Deploy to your private or public cloud instance on Azure, AWS, and Google Cloud Platform, or deploy in a private cloud or at the edge.

Proven Results in 6-Month Pilot **Visit C3.ai/get-started**