C3.ai

Enterprise AI Buyer's Guide

A comprehensive guide and checklist of requirements for evaluating and selecting Enterprise AI software

Table of Contents

How to Use This Guide

According to leading analysts and researchers, adoption of AI is a top enterprise IT priority and initiative:

- A June 2020 report from Gartner states that "smarter, faster, more responsible AI" is the #1 trend in data and analytics technology, as 75% of organizations will "shift from piloting to operationalizing artificial intelligence"
- An October 2019 survey by Forrester Research found that 53% of global data and analytics decision makers are in some stage of implementing artificial intelligence
- In March 2020, McKinsey reported that 76% of high-performing organizations have taken a standardized approach to AI technology, compared to only 18% of others

This guide is intended for business and IT professionals who need to evaluate and select Enterprise AI technology solutions for their organization. It is particularly designed for those who are looking to deploy Enterprise AI at large scale, as a strategic opportunity to drive digital transformation of their organizations. For these buyers, Enterprise AI technology represents a strategic, long-term investment in capabilities enabling their organizations to efficiently develop, deploy, and maintain multiple Enterprise AI applications across their business operations.

Given the current state of peak hype about AI, the marketplace for Enterprise AI technology has become noisy and confusing. Hundreds of vendors are attempting to position their offerings as AI solutions. This situation creates increasing difficulties for prospective buyers in (1) understanding what the specific requirements for Enterprise AI technology solutions are, (2) determining which vendors and products actually deliver against those requirements, and (3) discerning the differences among various offerings.

This guide provides a comprehensive, organized framework for cutting through the hype and evaluating offerings using specific, detailed criteria. These criteria are based on C3.ai's decade of experience in working with industry-leading companies to drive digital transformation through large-scale Enterprise AI deployments across multiple use cases and business processes. For example, one of the world's largest energy utilities based in Europe, is executing an Enterprise AI roadmap that addresses more than 130 use cases, ranging from predictive maintenance of the power grid to AI-enabled commodities trading. The company's use of AI to manage and optimize its 2-million-kilometer grid, comprising more than 40 million smart meters, is one of the world's largest industrial implementations of Enterprise AI. The value of its AI-driven digital transformation efforts is estimated at several hundred million euros annually.

At the heart of this guide is the Enterprise AI Requirements Checklist, that specifies all the capabilities an Enterprise AI technology platform must provide to deliver a complete, robust solution. Readers can use the Checklist to score and evaluate various offerings, applying their own weights to the criteria as dictated by their specific needs.

We recognize that a review of technical requirements is just one input into the overall evaluation process. This guide also includes a discussion of additional business requirements about a vendor's capabilities that buyers must also consider. The totality of assessing technical and business requirements, along with a proof of technology evaluation, will ultimately drive the final selection.

 C3.ai

What Is Enterprise AI?

Artificial intelligence is defined in various ways by analysts, practitioners, and observers – for example:

- "The ability of a machine to perform cognitive functions we associate with human minds, such as perceiving, reasoning, learning, and problem solving" (McKinsey & Company)
- "Advanced analysis and logic-based techniques, including machine learning, to interpret events, support and automate decisions, and take actions" (Gartner)
- "A branch of computer science dealing with the simulation of intelligent behavior in computers" (Merriam-Webster)

Enterprise AI is a powerful new category of software that is a core enabler of digital transformation. In the coming years, virtually every large organization will deploy dozens – often hundreds, or even thousands – of AI-enabled software applications across every aspect of their operations. These applications leverage 21st century technologies – including elastic cloud computing, big data, the internet of things, and advanced methods of artificial intelligence – to address a broad and growing range of use cases. A sampling of organizations that C3.ai works with indicates how AI is being applied to drive significant value:

- A major global oil and gas company has more than 280 AI initiatives in flight globally across its upstream, midstream, and downstream operations – a core component of the company's digital transformation strategy driving $1 billion in value through reduced costs, increased production, and higher margins.
- A leading multinational manufacturer is applying AI to optimize its entire global supply chain that spans more than 200 locations, enabling the manufacture and distribution of more than 40,000 different SKUs. AI-driven inventory optimization has helped improve safety stock levels by 10-20% and free up machine capacity by 3-5%.
- The US Department of Defense is deploying Enterprise AI across thousands of aircraft in multiple fleets – including E-3 Sentry, F-16, and F-35 – to predict aircraft subsystem failures before they occur, resulting in up to a 40% reduction in unscheduled maintenance and 6% increase in mission readiness.
- At an East Coast utility that powers a major US metropolis, its Enterprise AI platform provides mission-critical capabilities – continuously ingesting and analyzing data from more than 5 million smart meters to ensure the operational health of its network.

Despite the potential of AI to drive significant value, many organizations are failing in their digital transformation efforts. According to McKinsey & Company, fewer than 20% of enterprises achieve sustained value from their transformation initiatives. In contrast, the success rate of C3.ai customers approaches 100%. Common across these successful experiences is the adoption of an Enterprise AI platform.

This approach to Enterprise AI – unlike AI side projects that rarely progress beyond prototypes – is characterized by:

- **Scale:** Organizations deploy Enterprise AI at large scale across their entire value chain.

- **Value:** Enterprise AI deployments are focused on high-value use cases that will drive significant economic, social, and/or environmental impact.
- **Repeatability:** Enterprise AI is the systematic and repeatable application of AI, leveraging a common set of enterprise capabilities and assets to address a broad range of use cases.

Why Organizations Require an Enterprise AI Platform

Enterprise AI represents a new computing and programming paradigm. Distinctive to this new paradigm is the need to (1) aggregate and unify very large amounts of disparate data from multiple sources and in different formats, and (2) apply sophisticated AI and machine learning methods to these data at massive scale. In order to efficiently build, deploy, and scale high-performance Enterprise AI applications in a rapid, consistent, and repeatable manner, organizations require an Enterprise AI platform.

This is similar to what historically has occurred in prior examples of emerging new computing paradigms, such as internet computing, cloud computing, and mobile computing. For each of these new paradigms, new development platforms also emerged. These platforms provide comprehensive, unified capabilities that enable organizational stakeholders – software architects, engineers, data scientists, analysts, and technical administrators – to efficiently leverage these new technologies, and address the recurring requirements for building, testing, deploying, and operating these new classes of applications.

A decade ago, organizations that experimented with developing AI applications had no choice but to build these applications using a wide variety of primarily open source components of varying quality, created by numerous contributors with varying levels of experience and programming capability. None of these components were designed to work with all the other components required by an AI application, so the application developers had to write extensive code to ensure the interoperability of all the components. This approach produced highly complex and brittle applications that proved difficult to scale and nearly impossible to maintain.

Today, organizations are now able to build Enterprise AI applications repeatedly and at scale using a platform approach that removes the complexity and greatly simplifies the development and deployment of these applications.

A US-based multinational manufacturer with over $110 billion in revenue, for example, adopted an Enterprise AI platform approach after several years of efforts that had achieved some progress but not at the pace the company wanted. With its Enterprise AI platform, the company is addressing use cases ranging from predictive maintenance for paper manufacturing equipment to production scheduling optimization for petrochemical processing.

Similarly, one of the world's top 5 oil and gas companies had attempted to apply AI to predictive maintenance for hundreds of thousands of valves across multiple plants with limited success. By adopting an Enterprise AI platform, the company is now able to operate approximately 2 million AI machine learning models, enabling predictive maintenance for more than 500,000 valves. In just a single incident at one of its refineries, the use of AI saved $2 million by avoiding a potential valve failure.

Specifying the requirements of such a platform is the focus of this Enterprise AI Buyer's Guide.

Enterprise AI Market Landscape

The AI hype cycle has reached a peak as hundreds of vendors attempt to position their offerings as Enterprise AI solutions. Organizations are struggling to cut through the noise and understand the capabilities from various vendors and how products are differentiated. The market offerings being promoted as AI solutions fall into three basic categories:

Open Source Components

Open source software components – such as Hadoop, Cassandra, TensorFlow, Spark, and many others – are frequently positioned by their promoters as AI solutions. Each of these provides a valuable utility needed to develop and deploy Enterprise AI problem, but each addresses only a narrow portion of the total capabilities required. Using this approach, organizations must effectively create their own Enterprise AI platform by assembling, integrating, and maintaining dozens of such components. This approach has proven ineffective for developing and operating Enterprise AI applications at scale in an efficient, repeatable, and predictable manner.

Cloud Service Providers

Leading cloud service providers (CSPs) such as AWS and Azure provide an increasingly large array of native services that can be used to assemble AI applications. The advantage of this approach over attempting to build AI applications using open source components is that a CSP's native services have been designed to work on the CSP's infrastructure. However, these services often are not tightly integrated with one another and therefore require extensive coding – as well as considerable expertise and experience with the CSP's infrastructure – to make them all work together in a functioning application.

Moreover, each CSP's native services are tightly bound to the CSP's infrastructure, and are not portable from one CSP to another. So, an application developed on AWS using AWS native services, for example, would have to be substantially re-written to run on Azure. In addition, this approach precludes the ability to build an application using services from different CSPs (for instance, an application using AWS's image recognition service in conjunction with Google's geospatial capabilities).

Enterprise AI Platforms

An Enterprise AI Platform is designed to provide a cohesive set of capabilities in a unified, pre-integrated suite to build, deploy, and operate Enterprise AI applications. The intention of a platform is to simplify and accelerate the development and deployment of Enterprise AI applications. This buyer's guide focuses on the Enterprise AI Platform category of offerings.

10 Core Principles of Enterprise AI

Transformative impact from Enterprise AI comes from applying AI at scale across an organization's entire value chain. Significant opportunity for innovation and competitive advantage lies in applying AI to re-think how businesses operate and deliver dramatic improvements in how companies engage customers, make better use of their workforce, and improve business operations. The use cases for AI in banking, for instance, are numerous. AI applied to data produced by transaction and order systems, product systems, client masters, and document repositories can proactively identify the need to address corporate cash churn or to prioritize anti-money laundering efforts. AI and optimization techniques can be used to anticipate fluctuations in customer demand or supply disruptions, to better inform securities lending efforts, or for early identification of loan application risk.

Representative Value Chain: Banking

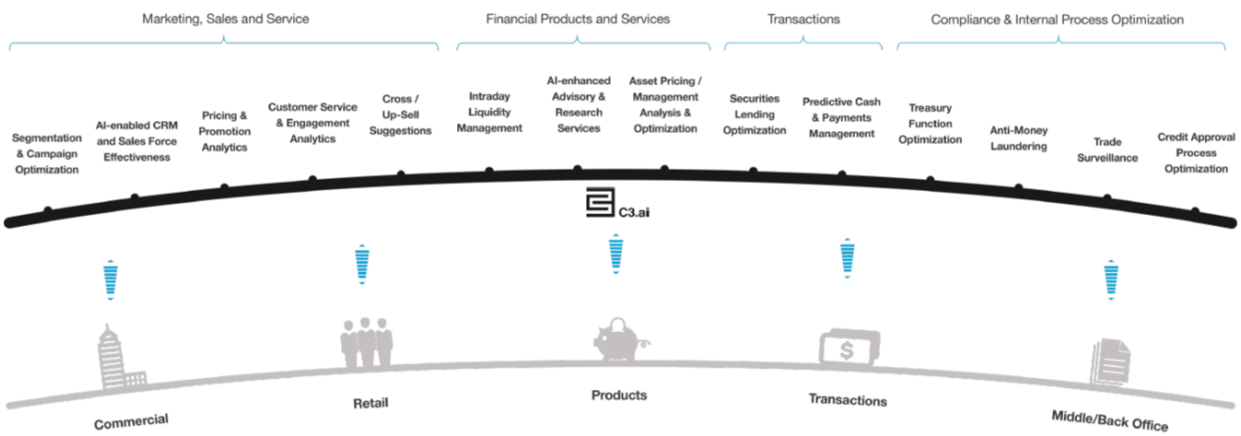

The information technology challenges in delivering these AI applications at scale across the enterprise are daunting. Based on C3.ai's decade of experience helping global organizations apply Enterprise AI across multiple industries – including manufacturing, aerospace, oil and gas, defense, healthcare, and utilities – we have identified and codified 10 core capabilities for a complete Enterprise AI platform, described below. In a following section, we have organized the specific requirements for each of the 10 core capabilities in a convenient, comprehensive Enterprise AI Requirements Checklist.

1. Data Aggregation: Unified Federated Data Image Across the Business

Process re-engineering across a company's business requires integrating data from numerous systems and sensors into a unified federated data image, and keeping that data image current in near real-time as data changes occur. The baseline capability required is aggregation and processing of rapidly growing petabyte-scale datasets continuously harvested from thousands of disparate legacy IT systems, internet sources, and multi-million sensor networks. In the case of one Fortune 500 manufacturer, for example, the magnitude of the data aggregation problem is 50 petabytes fragmented across 5,000 systems representing customer, dealer, claims, ordering, pricing, product design, engineering, planning, manufacturing, control systems, accounting, human resources, logistics, and supplier systems fragmented by mergers and acquisitions, product lines, geographies and customer engagement channels (i.e., online, stores, call center, field).

 |

To facilitate data integration and correlation across these systems requires a Data Integration service with a scalable enterprise message bus. The Data Integration service should provide extensible industry-specific data exchange models, such as HL7 for healthcare, eTOM for telecommunications, CIM for power utilities, PRODML and WITSML for oil and gas, and SWIFT for banking. Mapping source data systems to a common data exchange model significantly reduces the number of system interfaces required to be developed and maintained across systems. As a result, deployments with integrations to 15 to 20 source systems using an Enterprise AI platform with a Data Integration service will typically take 3–6 months as opposed to years.

2. Multi-Cloud Computing and Data Persistence

Cost effectively processing and persisting large-scale datasets requires an elastic cloud scale-out/in architecture, with support for private cloud, public cloud, or hybrid cloud deployments. Cloud portability is achieved through container technology (for example, Mesosphere). An Enterprise AI platform must be optimized to take advantage of differentiated services. For example, the platform should enable an application to take advantage of AWS Kinesis when running on AWS and of Azure Streams when running on Azure.

The platform must also support Multi-Cloud operation. For example, the platform should be able to operate on AWS and invoke Google Translate or speech recognition services and access data stored on a private cloud. It should also be possible for an instance of the platform to be deployed in-country – for example, on Azure Stack – so that it conforms to data sovereignty regulations.

The platform needs to support installation in a customer's virtual private cloud account (e.g., Azure or AWS account) and support deployment in specialized clouds such as AWS GovCloud or C2S with industry- or government-specific security certifications.

Data persistence of the unified data image requires a multiplicity of data stores depending on the data and anticipated access patterns. Relational databases are required to support transactions and complex queries, and key-value stores for data such as telemetry requiring low-latency reads and writes. Other stores, including distributed file systems such as HDFS, are required for support of unstructured audio or video, multi-dimensional stores, and graph stores. If a Data Lake already exists, the platform can map and access data in-place from that source system.

3. Edge Computing

To support low-latency compute requirements or situations where network bandwidth is constrained or intermittent (e.g., aircraft), an Enterprise AI platform must enable local processing and the ability to run AI analytics, predictions, and inferences on remote gateways and edge devices.

4. Platform Services and Data Virtualization for Accessing Data In-Place

AI applications require a comprehensive set of platform services for processing data in batches, microbatches, real-time streams, and iteratively in memory across a cluster of servers to support data scientists testing analytic features and algorithms against production-scale data sets. Secure data handling is also required to ensure data is encrypted while in motion or at rest. The Enterprise AI platform architecture should allow pluggability of these services without the need to alter application code.

The architecture should also support data virtualization, allowing application developers to manipulate data without knowledge of the underlying data stores. The platform needs to support database technologies including relational data stores, distributed file systems, key-value stores, graph stores as well as legacy applications and systems such as SAP, OSIsoft PI, and SCADA systems.

5. Enterprise Semantic Model

Re-engineering processes across an organization requires a consistent semantic (object) model across the enterprise. An Enterprise AI platform must support a semantic model that represents objects and their relationships independent of the underlying persistence data formats and stores. In contrast to passive entity / object models in typical modeling tools, the object model must be active and interpreted by the platform at runtime providing significant flexibility to handle object model and schema changes. The platform should support changes to the object model, versioned, and immediately active without need to re-write application code.

Enterprise Semantic Model

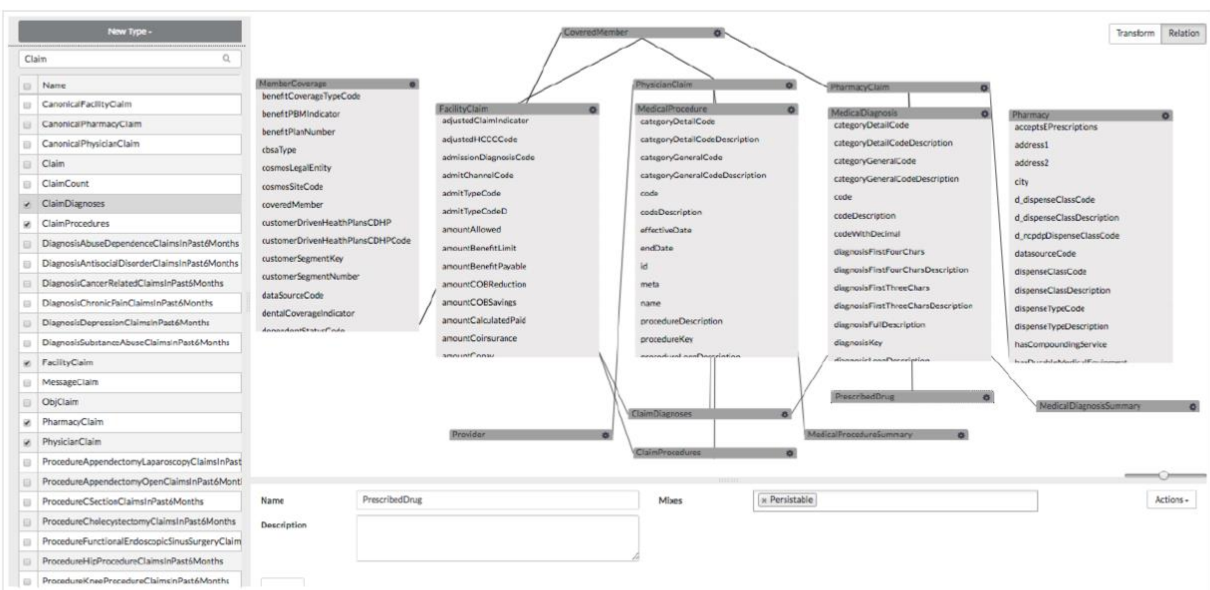

6. Enterprise Microservices

In order to enable developers to rapidly build applications that leverage the best components, an Enterprise AI platform must provide a comprehensive catalog of AI-based software services. This catalog of AI microservices should be published and available enterprise-wide, subject to security and authorization access controls.

7. Enterprise Data Governance and Security

The Enterprise AI platform must provide robust encryption, multi-level user access authentication, and authorization controls. Access to all data objects, methods, aggregate services, and ML algorithms should be subject to authorization. Authorization should be dynamic and programmatically settable; for example, authorization to access data or invoke a method might be subject to the user's ability to access specific data rows. The platform must also provide support for external security authorization services – for example, centralized consent management services in financial services and healthcare.

8. System Simulation Using AI and Dynamic Optimization Algorithms

An Enterprise AI platform must support an integrated full life-cycle algorithm development experience so that data scientists can rapidly design, develop, test, and deploy machine learning and deep learning algorithms. Data scientists should be able to use the programming language of their choice – Python, R, Scala, Java – to develop and test machine learning and deep learning algorithms against a current production snapshot of all available data. This ensures that data scientists can achieve the highest levels of machine learning accuracy (precision and recall).

The platform must enable machine learning algorithms to be deployed in production without the effort, time, and errors introduced by translation to a different programming language. Machine learning algorithms should provide APIs to programmatically trigger predictions and re-training as necessary. AI predictions should be conditionally triggered based on the arrival of dependent data. AI predictions can trigger events and notifications or be inputs to other routines including simulations involving constraint programming.

9. Open Platform

The ability to interoperate with other technologies, products, and components is essential to maximizing developer and data science productivity, enabling collaboration, and keeping pace with ongoing innovation. An Enterprise AI platform must be open, providing plug-ins and flexibility for data scientists and developers – including IDEs and tools, programming languages, DevOps capabilities, and others.

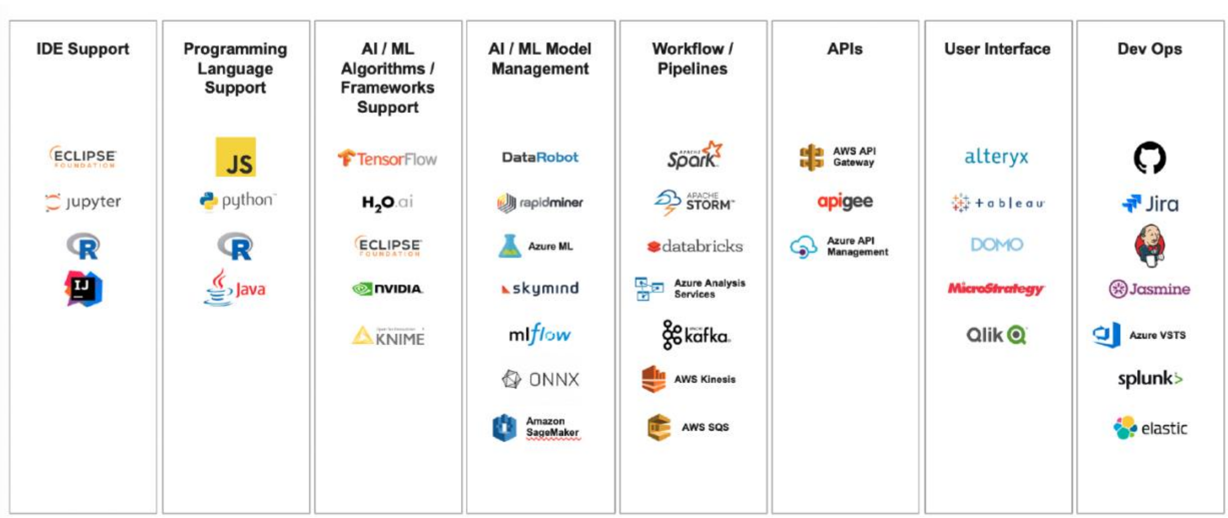

The platform must support standards-based interfaces (APIs), open source machine learning and deep learning libraries, and third-party data visualization tools. The platform must enable the incorporation of any new open source or proprietary software innovations without adversely affecting the functionality or performance of an organization's existing applications.

10. Common Platform for Collaborative Development between Software Developers and Data Scientists

Data scientists typically work in isolation, developing and testing machine learning algorithms against small subsets of data provided by IT from one or more disparate source systems. The bulk of their time is spent on data cleansing and data normalization to represent the same entities, measures (units), states (e.g., status codes), and events consistently in time and across systems, and to correlate ("join") data across systems. The resulting algorithms, typically written in Python or R, do not typically conform to IT standards and may require rewriting to a different programming language such as Java. Furthermore, the efficacy of the algorithm is almost always sub-optimal since it has not been tuned against a representative production data set.

To overcome these obstacles, an Enterprise AI platform must allow data scientists to develop, test, and tune algorithms in the programming language of their choice against a snapshot of all available production data. To accelerate development, the platform must enable data scientists to leverage work completed on the platform by data engineers and application developers to handle data cleansing, data normalization, object modeling, and representation. And it must provide microservices to focus on analytic feature development for classic machine learning or deep learning models. The resulting machine learning algorithm should be immediately deployable in production and available as a microservice through a standard RESTful API.

All data objects, methods, aggregate services, and ML algorithms should be accessible through standard programming languages (R, Java, JavaScript, Python) and IDEs (Eclipse, Azure Developer Tools). The Enterprise AI platform must provide a complete, easy-to-use set of visual tools to rapidly configure applications by extending the metadata repository. Metadata repository APIs are also required to synchronize the object definitions and relationships with external repositories or for introspection of available data objects, methods, aggregate services, and ML algorithms. Application version control must be available through synchronization with common source code repositories such as git.

Model-Driven Architecture

The architecture requirements for an Enterprise AI platform are uniquely addressed through a Model-Driven Architecture. This architecture abstracts application and machine learning code from the underlying platform services and provides a domain-specific language (annotations) to support highly declarative, low-code application development.

The model-driven approach provides an abstraction from the underlying technical services (for example, queuing services, streaming services, ETL services, data encryption, data persistence, authorization, authentication) and simplifies the programming interface required to develop AI applications to a type system interface.

The model is used to represent all layers of an application including the data interchange with source systems, application objects and their methods, AI-machine learning algorithms and the application user interface. Each of these layers are also accessible as microservices.

Model-Driven AI Architecture

An example of a proven model-driven AI architecture to support Enterprise AI applications is the C3 AI SuiteTM. The C3 AI Suite allows small teams of 5 to 10 application developers and data scientists to collaboratively develop, test, and deploy large-scale production AI applications in one to three months. The platform is proven in 30 large-scale deployments across industries including Energy, Manufacturing, Aerospace / Defense, Healthcare, and Financial Services. A representative large-scale C3 AI Suite deployment processes AI inferences at a rate of a million messages per second against a petabyte-sized unified federated cloud data image aggregated from 15 disparate corporate systems and a 40-million sensor network. Global 1000 organizations have successfully used the C3 AI Suite as an Enterprise AI platform to deploy full-scale production deployments in 6 months and enterprise-wide digital transformations with over 20 AI applications in 24 to 36-month timeframes.

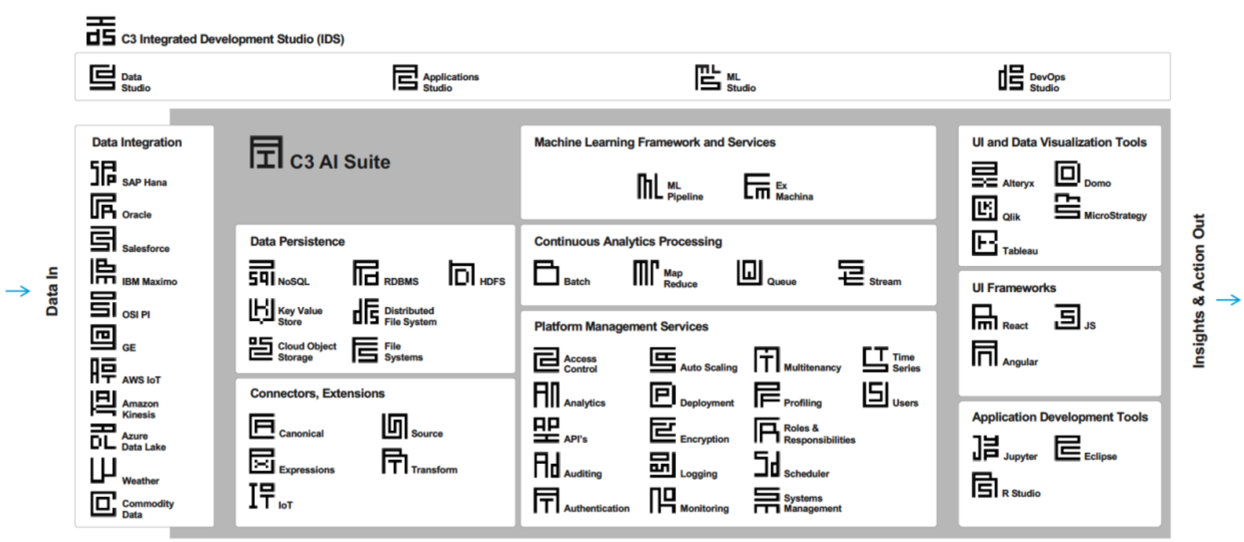

Type System

An example of a type system in a model-driven architecture and scalable Enterprise AI platform is the C3.ai Type System™. The C3.ai Type System is a data object-centric abstraction layer that binds the various C3 AI Suite components, including infrastructure and services. It is both sufficient and necessary for developing and operating complex predictive analytics applications in the cloud.

The C3.ai Type System is the medium through which application developers and data scientists access the C3 AI Suite, C3.ai Data Lake, C3.ai Applications, and applications and microservices. Examples of C3.ai Types include data objects (e.g., customer, product, supplier, contract, or sales opportunity) and their methods, application logic, and machine learning classifiers.

The C3.ai Type System allows programs, algorithms, and data structures – written in different programming languages, with different computational models, making different assumptions about the underlying infrastructure – to interoperate without knowledge of the underlying physical data models, data federation and storage models, interrelationships, dependencies, or the bindings between the various structural platform or cloud infrastructure services and components (e.g., RDBMS, NoSQL, ETL, SPARK, Kafka, SQS, Kinesis, object models, classifiers, data science tools, etc.). The C3.ai Type System provides RESTful interfaces and programming language bindings to all underlying data and functionality.

Leveraging the C3.ai Type System, application developers and data scientists can focus on delivering immediate value, without the need to learn, integrate, or understand the complexities of the underlying systems. The C3.ai Type System enables programmers and data scientists to develop and deploy production AI, big data, and predictive analytics applications in one-tenth the time at one-tenth the cost of alternative technologies.

To improve manageability, Types support multiple object inheritance (allowing objects to inherit characteristics from one or more other objects). For example, a building might have characteristics of both a residential and commercial use building.

The Type system, through inherent dataflow capabilities, automatically triggers the appropriate processing of data changes by tracing implicit dependencies between objects, aggregates, analytic features and machine learning algorithms in a directed acyclic graph.

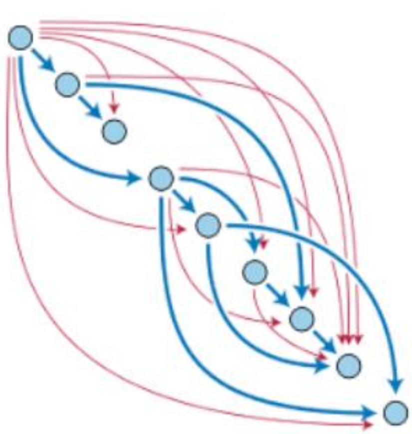

Figure: Data changes automatically trigger data pipelines including re-computation of analytic features, aggregates and AI predictions

The C3.ai Type System is accessible through multiple programming language bindings (i.e., Java, JavaScript, Python, Scala, and R), and Types are automatically accessible through RESTful interfaces allowing interoperability with external systems.

Modeling AI Applications

C3.ai applications are expressed through the following model artifacts:

1. Metadata – example - used to declare objects and their attributes
2. Annotations – example - used to indicate the underlying data store to persist an object's data
3. Expressions – example - used to represent functions which perform mathematical operations on input data without programming logic
4. Programming Logic – example - used when necessary to express logic beyond metadata, annotations and expressions -

These artifacts are used to represent the following interdependent "layers" of an application:

1. Data Exchange Object Model, Mappings, and Transformations – metadata representing a data model optimized/denormalized for data interchange between systems, mappings and data transformations to target an application object model.
2. Application Object Model – metadata representing a data model optimized for development of applications and machine learning models.
3. Methods – business logic representing actions (e.g., enroll a customer, maintain equipment) on objects expressed in code
4. Analytic Features – expressions which perform spatial and/or temporal mathematical operations on input data
5. AI Machine Learning Model – combination of analytic features and AI-ML algorithm (e.g., NLP pipeline, logistic regression, TensorFlow…)
6. Security Services – annotations and expressions applied to Objects, Methods, Features and Machine Learning Models.
7. User Interface (Optional) – metadata representing user interface fields, their bindings to objects and object properties, and user interface objects (charts, grids, etc.,)

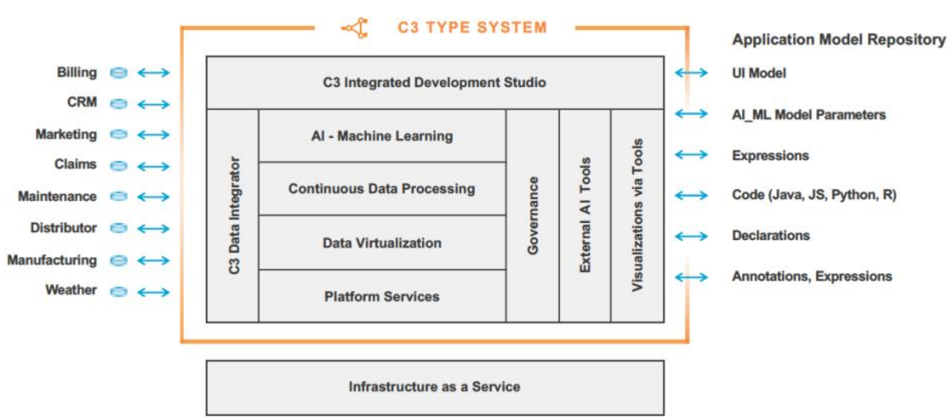

Model artifacts are extensible and upgradeable:

Model Extensibility

- Meta Data Extensions
- Code Extensions

Model Upgradeability

- Platform Upgrades
- Application Upgrades
- Preserving Customer Extensions

The benefits of a model-driven architecture include: 1) future-proofing investments in application and microservices code as underlying infrastructure services rapidly evolve; 2) less application code is required to be written, quality assured, and maintained, resulting in significantly faster application development and lower total cost of ownership; and 3) upgradeability.

Enterprise AI Requirements Checklist

By applying your own weights to this comprehensive requirements list as directed by your organization's specific needs, you can use this Checklist to score offerings under consideration.

Weight	Does the solution offer:	Vendor 1	Vendor 2	Vendor 3
DATA AGGREGATION				
	Multi-Dimensional Data Services			
	Database Services • Relational • Key-Value Store • Data Auditing • Data Versioning • Pre-built Utility Data Model • Data Validation			
	Canonical-based Data Integration			
	REST API Enabled Datasets			
	In-Memory Distributed Data Services			
	Data Virtualization (Distributed Query Framework) • Relational Databases: Oracle, SAP Hana, Postgres, RDS, SQL Server • Distributed File Systems: HDFS, S3,… • Application Systems: SAP, Oracle, Maximo • Key-Value Stores: Cassandra, HBase, DynamoDB, Azure Cosmos DB			
	Data Types: Structured & Unstructured (Image, Video, Sound…)			
MULTI-CLOUD COMPUTING				
	Multi-Cloud Deployment • Public Cloud (Azure, AWS, Google, etc.…) • Private Cloud			
EDGE COMPUTING				
	Supports local processing and the ability to run AI analytics, predictions, and inferences on remote gateways and edge devices			
PLATFORM SERVICES				
	Model Driven Architecture			
	Connectivity to Open Source Data Repositories			
	Automatic Hot / Cold Storage for High Volume Data			
	Single Deployment, Multi-Tenant Platform			
	Support Multiple Analytical Runtime Environments Concurrently			
	AI Model Operations & Monitoring			
	Large-Scale Production AI Model Deployments			
	Business Process Management: Workflow & Orchestration			
	Business Rules Engine			
	Event points for procedural logic			
	Continuous/Event Processing Services • Distributed Stream Computing • Complex Event Processing			
	Private PaaS • Virtual Private Instance • Appliance Datacenter Deployment Portals			
	OpenShift / Kubernetes Based Deployment			
	Deployment with Prescribed CI/CD process			
ENTERPRISE SEMANTIC MODEL				
	Unified, Extensible Type/Object System • Pre-built Utility Object Model • Object-Relational/Key-Value Mapping • Extensible Type System • Data Model, Object Model, Integration Model, Expressions, Analytic Features, User Interface, App Logic			
	Application Logic Encoded in Meta-Data			
	Declarative Representation of Application Semantics to Reduce Code • Data Exchange / Canonical Models • Data Objects and Data Inter-relationships			

	• Workflows • Analytics/AI ML • User Interface			
	Object Level Governance and Provenance			
	Object Inheritance and Multi-Inheritance to Support Variants			
ENTERPRISE MICROSERVICES				
	Analytics Services • Fine Grained, Continuous Analytical Processing • Analytic Metrics • Analytic Expressions/Features • Predictive Modeling • Machine Learning Classifiers			
	MapReduce Service			
	Data Science • NLP, Speech, Image, Facial • Libraries: Scikit learn, TensorFlow, Caffe, Torch,… • Supervised, Unsupervised • Classification, Regression			
	Data Discovery - Data and ML Model Catalog			
	No Code Data Discovery and Analysis			
	Jupyter Notebook Service			
	Auto ML / Hyperparameter Optimization			
	Advanced Time Series Analysis			
	Multi-Framework, Composable, Machine Learning Pipelines			
ENTERPRISE DATA SECURITY				
	Security Services • Authentication/Authorization • Roles/Responsibilities • Field, Row, Object Visibility • Data encryption at rest • Data encryption at rest			
	Security Certifications: SOC 2 Type 2; HITrust; FedRamp			
SYSTEM SIMULATION USING AI AND DYNAMIC OPTIMIZATION ALGORITHMS				
	Application Development Lifecycle Management • Integrated Graphical Design Environment • Provisioning • Upgrades • Monitoring • Auto-scaling			
OPEN PLATFORM				
	Open Architecture – Support for 3rd-Party Development & Data Science Tools			
	Data Integration & Ingestion Capabilities • ETL • Synchronous/Asynchronous Messaging • Integration Process • REST APIs • Utility Canonical Object Model • Canonical Object to Data Model Mapping			
	Open Source Frameworks for Analytics			
	Publish AI Models with REST API			
	BI Tools – Tableau, Qlik, MicroStrategy			
COMMON PLATFORM FOR COLLABORATIVE DEVELOPMENT				
	Integrated Graphical Designers • Visual User Interface Designer • Data Model Designer • Object Model Designer • Integration Designer • Analytics Designer • Visual Analytics Designer • Application Logic Configuration and Scripting • Data Explorer • Application Provisioning			
	UI / App Framework for ML Applications			
	Programming Languages: Java, JavaScript, R, Python, Scala			

	Mobile, Responsive design			

Additional Vendor Evaluation Criteria: Business Requirements

In addition to a full technical evaluation, buyers of Enterprise AI solutions should also evaluate vendors on the basis of the following business criteria.

Experience	Does the vendor have a substantial track record of market experience and success?
Leadership Team	Does the vendor have a leadership team committed to customer success, with a proven track record of delivering technology solutions at scale?
Referenceable Customers	Does the vendor have referenceable customers who have deployed its solutions at enterprise scale with significant measurable results?
Partner Ecosystem	Does the vendor have significant relationships with other relevant leading technology and services providers?
Methodology, Best Practices, and Knowledge Transfer	Does the vendor possess a proven methodology and best practices for customer success, and is there a robust process in place to transfer this knowledge and skills to your organization?

About C3.ai

C3.ai is a leading AI software provider for accelerating digital transformation. C3.ai delivers the C3 AI Suite™ for developing, deploying, and operating large-scale AI, predictive analytics, and IoT applications in addition to an increasingly broad portfolio of turn-key AI applications. The core of the C3.ai offering is a revolutionary, model-driven AI architecture that dramatically enhances data science and application development.

Proven Results in 8-12 Weeks | **Visit c3.ai/get-started**

 |