

C3 Code

Frequently Asked Questions

What is C3 Code?

C3 Code is the agentic app builder within the C3 Agentic AI Platform. Describe a business problem in natural language, and its agentic coding assistants autonomously design, build, test, and deploy a production-grade Enterprise AI application — no coding required. Work that once took a team of engineers weeks to months can now take one person hours to days. “Production-grade” means our apps integrate with enterprise data, applies machine learning, and ships with full security, access control, and governance. This is not assisted development; it is AI designing and building Enterprise AI.

Who is C3 Code for?

Business analysts, IT, developers, and data scientists. You don’t need to know how to code — you need to understand your business problem.

What can C3 Code do?

From a single natural-language prompt or product specification, C3 Code delivers a complete, deployed application:

- **Data integration** — connects to Snowflake, Databricks, Oracle, Salesforce, ERP systems, and more; eliminates weeks of manual ETL
- **Enterprise data modeling** — generates and extends C3 AI Types, the production-grade ontology of every C3 AI application
- **Machine learning** — configures proven domain algorithms (predictive maintenance, demand forecasting, anomaly detection, computer vision) without ML engineering resources
- **Autonomous agents** — configures agentic workflows that run business processes: optimizing inventory, updating forecasts, triaging equipment alerts
- **Full-stack UI** — dashboards, KPI tiles, geospatial maps, chat interfaces
- **One-click governed deployment** — security, access control, and audit trails built into the architecture, plus production monitoring

What types of applications, agents, and workflows can you build with C3 Code?

You can build any Enterprise AI use case — Data/BI applications, autonomous agents, multi-step agentic workflows, and generative AI apps. Example use cases include:

- **Supply chain & logistics** — inventory visibility and shortage detection, parts track-and-trace, demand forecasting, freight and route optimization, supplier onboarding and risk
- **Asset performance & reliability** — predictive maintenance, anomaly detection, root-cause analysis, maintenance (CMMS) and spare-parts optimization
- **Manufacturing & operations** — production scheduling and sequencing, throughput and quality (SPC) monitoring, energy and emissions optimization, shift and calendar sync

C3 Code: Frequently Asked Questions

- **Finance & back-office** — EBITDA/revenue forecasting and scenario analysis, AP/invoice automation, order-to-cash, financial close, ESG reporting
- **Customer, sales & service** — customer 360 and churn, pipeline forecasting, quoting/CPQ, customer service agents, field-service dispatch
- **Risk, compliance, IT & security** — contract review and lifecycle, regulatory-change monitoring, AML investigation, third-party risk, ticket and alert triage, incident response

What can C3 Code not do (yet)?

Extending existing C3 AI applications — such as *C3 AI Reliability*, *C3 AI Production Schedule Optimization*, and *C3 AI Inventory Optimization* — is on our roadmap and will arrive in future releases. C3 Code is purpose-built for Enterprise AI on the C3 Agentic AI Platform, so it is not the tool for: developing custom code inside third-party enterprise systems (e.g., SAP ABAP customizations); provisioning custom infrastructure or standalone web stacks (e.g., new Kubernetes clusters, Next.js or Vercel apps); or native mobile apps (e.g., Swift/iOS or Android) — though it does build responsive UIs that work well on mobile.

What LLMs does C3 Code support?

C3 Code is LLM-agnostic — it works with Claude (Anthropic), Codex/GPT (OpenAI), Gemini (Google), Grok (xAI), Mistral, Llama (Meta), or any other capable model, and you can switch without re-architecting the application. Claude is the default model.

How is C3 Code different from other coding agents (Cursor, GitHub Copilot, Claude Code, Codex)?

Those are excellent general-purpose tools for writing code faster; C3 Code autonomously builds, deploys, and operates production-grade Enterprise AI applications. It delivers every capability of a leading coding agent and adds the enterprise layer only C3 AI provides: prebuilt domain intelligence, native data and ML infrastructure, governance, and production operations. General-purpose tools start from a blank canvas on every project; C3 Code extends proven, production-hardened assets.

Capability	General Purpose Coding Agent	C3 Code
Agentic code generation, Python/JS, IDE & Git	—	✓
LLM-agnostic (Claude, Codex/GPT, Gemini, Grok, Llama...)	—	✓
End-to-end autonomous, full-stack app generation	—	✓
C3 AI Type System, ontology & domain-specific language	—	✓
40+ prebuilt app packages & production domain AI algorithms	—	✓
Native data integration, virtualization & ML pipelines	—	✓
One-click governed deploy, RBAC, audit trails & monitoring	—	✓

Is C3 Code more cost effective than other coding agents?

Yes, on two fronts. It saves **developer time** because every enterprise-grade application primitive — data integration, the ontology and C3 AI Type System, ML pipelines, agents, governance, and deployment — is already prebuilt in the C3 Agentic AI Platform, so teams configure and extend rather than build from scratch. And our benchmarks indicate C3 Code is at least **50% more token efficient** than other coding agents for an equivalent application: it is LLM-agnostic (each task runs on the most cost-effective model), optimizes token and prompt caching (reusing platform context instead of re-sending it), and its prebuilt packages and algorithms mean the model generates far less from scratch.

How is C3 Code different from frontier models?

Frontier models can generate code and prototype applications. What they cannot do is deliver a governed, enterprise-scale AI application — one that integrates with your data infrastructure, runs validated domain algorithms, enforces role-based access control, and deploys through secure pipelines on industry-specific data models refined over years. Enterprise AI is not a coding problem. It is a systems problem, a data problem, a governance problem, and a domain knowledge problem. C3 Code solves all of them simultaneously. Better models make C3 Code stronger, not weaker. C3 Code is model-agnostic — frontier advances flow directly into the platform. But what any model gains through C3 Code is something no general-purpose environment can replicate: native intelligence of the C3 Agentic AI Platform — the C3 AI Type System, the domain-specific language, 40+ Enterprise AI application packages, domain AI algorithms, and the full architectural corpus from a decade of enterprise deployments. The moat is not the model; it is the enterprise intelligence layer.

Does C3 Code change the scale of what a single person can deliver?

Yes. C3 Code changes how value gets delivered in Enterprise AI, not just for C3 AI, but for the entire industry. C3 Code dramatically expands who can build Enterprise AI. The focus now is getting C3 Code into the hands of customers and partners and expanding the addressable market.

TRUST, IP & GOVERNANCE

How do I trust what C3 Code builds, and do I still need developers?

C3 Code has guardrails built into the workflow, so you ship a working, reviewed application rather than unverified output. Generated code is grounded in the C3 AI Type System and governed data — constrained by the platform's schema and validation rather than free-form, which sharply limits room for hallucinated logic. It produces type-safe data models, APIs, agents, ML, and UI, then writes tests and opens pull requests. Humans stay in control: engineers review the generated code and tests, catch and correct issues, make last-mile changes, and approve the governed deployment. The result is production-grade from the first pass, with testing and human review catching potential problems before go-live.

Can developers still access and modify the generated code?

Yes. A full IDE is available so developers can review and change any generated code, run tests, and work through their existing Git workflow — C3 Code fits into the tools your team already uses.

Who owns the applications and IP that C3 Code generates?

If you build it, you own it — applications and IP your team creates with C3 Code belong to you. When C3 AI employees or contractors build an application on your behalf, C3 AI owns that IP. Specifics are set out in your commercial agreement.

How does C3 Code handle security, governance, and my data?

Applications inherit the platform's security model — role-based access control, data governance, encryption, audit trails, and agent evaluations are architectural, not add-ons — with production monitoring of data ops, jobs, queue health, and agent ops built in. The agent operates within each user's existing permissions, and C3 AI does not use your data to train models. C3 AI's federal heritage supports the most demanding regulated and defense deployments.

Is there vendor lock-in with C3 Code?

C3 Code is open by design — no model lock-in, no toolchain lock-in. Use Claude, Codex/GPT, Gemini, Grok, Mistral, Llama, or any capable model. Applications deploy on your cloud — AWS, Azure, or Google Cloud.

DEPLOYMENT & ARCHITECTURE

Can C3 Code run in my cloud, on-prem, or air-gapped environment?

Yes. C3 Code and the applications it builds can run in your cloud (AWS, Azure, or Google Cloud), on-premises, or in isolated / air-gapped environments to meet security and regulatory requirements. Customer-hosted and air-gapped deployments carry additional charges.

What data sources does it connect to out-of-the box?

C3 Code builds on C3 AI Data Fusion, so it connects to your systems out of the box — 130+ prebuilt enterprise connectors to common databases, data warehouses, and applications (e.g., Snowflake, Databricks, Oracle, Salesforce, SAP, and ERP/MRP systems) across cloud, on-prem, and streaming sources (Kafka, Kinesis, Event Hub). It ingests structured, unstructured, and sensor/time-series data, maps it to C3 AI canonical models, and can virtualize sources to avoid duplicative ETL. Custom connectors can be authored when a source isn't already covered.

BUYING & OPERATING

What is C3 Code priced?

C3 Code is part of the C3 Agentic AI Platform, which is priced on usage- and seat-based pricing.

- \$0.55/vCPU-hour
- \$2.50–\$9.50/vGPU-hour by GPU class.
- Adding a C3 AI application module adds a 30% surcharge on that module's vCPU/vGPU usage fee.
- First 20 developers are included. Additional developers at \$200/developer/month.
- Term commitments result in a discount on the list price (on usage only, not seats):
 - 1 year → 15% off
 - 2 years → 25% off
 - 3 years → 30% off
- C3 AI does not charge any token fees of its own. For C3-hosted deployments: cloud hosting and LLM provider token fees pass through to the customer at cost.
- Center of Excellence support and customer-hosted deployments are priced separately.

How is usage metered, and what controls do I have to manage or reduce usage?

Usage is metered in vCPU/vGPU-hours, covering compute consumed by C3 Code and deployed applications. If the coding agent calls a C3 managed LLM provider, that usage is metered separately in tokens and passed through at cost.

You have several controls to manage or reduce usage:

- **Visibility and caps** — usage analytics with real-time usage tracking, plus per-user token, workspace, and deployment quotas
- **Model governance** — restrict the LLM models that are available to your users
- **Compute efficiency** — hibernate environments not needed continuously, set autoscaling min/max thresholds for deployed applications

What happens at the end of the Accelerator — can I export what I built?

At the end of the Accelerator, convert to month-to-month usage- and seat-based pricing to keep building and running — nothing is torn down. Users can export the code. Applications built with C3 Code are native C3 Agentic AI Platform applications and can only operate as part of the platform; they are not designed to be exported to run standalone outside it. While on the platform, developers retain full access to review and modify the underlying code in the IDE.

What support, training, and onboarding does C3 AI provide?

C3 Code ships with full documentation, and its agentic assistant has the entire C3 AI knowledge corpus built in to answer build-related questions as you work. C3 AI's Accelerator program trains your team to build Enterprise AI applications and stands up your first use cases side-by-side (see "How do I get started?"). Center of Excellence support is available for ongoing enablement, priced separately.

Do I need C3 AI services, or can my team build independently?

Your team can build applications independently with C3 Code. The C3 AI Center of Excellence is available if you want help building more complex, multi-system applications, at additional cost.

How do I get started?

1. **See it live** — schedule a demo at c3.ai/contact-us and watch C3 Code build a production application in front of you.
2. **Sign up for the C3 Code Accelerator** — C3 AI builds your first application and use cases side-by-side with your team.
3. **Keep building** — When you're ready, convert to seat + usage-based pricing, with usage rate discounts for multi-year commitments.

Start Building.

See C3 Code build a production application live. Visit C3.ai/get-started