

FIRST CALL DECK
May 2026

BUILT BY
DinoCloud × AWS

INDUSTRY
Retail · CPG · Logistics

a productized aws-native solution

Troodon.

The AI commercial analyst that scales the best of your team —
every customer, every category, every week.

a creation of



today

Today's outline.

A 30-minute walkthrough of the product, the data foundation it sits on, the customer reference, and how we move forward.

01 DinoCloud at a glance

02 Production AI in LATAM retail

03 The diagnostic doesn't scale

04 Meet Troodon

05 Architecture & AWS stack

06 Approach & methodology

07 Molinos Río de la Plata — in production

08 Engagement model & next steps

about dinocloud

An AWS Premier partner — built for enterprise customers.

10 years building production AWS workloads for regulated and data-intensive industries. Specialized practice in Generative AI, Data & Analytics. Delivery hubs in Argentina and the U.S.

Premie

AWS Tier Services Partner
since 2024

10+

Years building on AWS in
LATAM and the U.S.

7

Active AWS competencies —
AI, Migration, Security,
DevOps, MSP, CloudOps,
FinServ

AR • US

Delivery footprint across the
Americas

industry track record

Production-ready AI for Retail, CPG & logistics.

Selected references where DinoCloud delivered AWS-native data and AI workloads in retail-adjacent verticals.

CPG · Food manufacturing

Molinos Río de la Plata

Troodon — flagship case

Pastas Secas pilot with Carrefour (KA) and Vemart (distributor). Bedrock Knowledge Bases + Text-to-SQL on Redshift. SAP → Data Lake in 8 weeks.

Logistics · Largest private courier (AR)

Andreani

Cloud platform engineering on AWS. Data foundation for operational analytics. Mission-critical workloads in production today.

Retail · Consumer electronics

Naldo

AWS migration and modernization. Data and analytics platform supporting e-commerce and store operations workloads.

the problem

Commercial teams analyze the few accounts they can.

The remaining 80%+ of accounts and SKUs run on intuition, monthly Excel exports, and the same playbook for every customer.

Data silos

Scentia, SAP and distributor data live in three different systems.

Sell-out, sell-in, distribution and pricing have no shared semantic layer — analysts stitch them by hand.

Coverage

Talent doesn't scale to 47 categories × every customer.

Categories, brands, SKUs, distributors and KAs multiply combinatorially.
Only Tier-1 accounts get the full diagnostic.

Time lag

Insights arrive ~30 days late.

Market data has a one-month lag and CIC is monthly granular. By the time a gap is found, the cycle is over.

Drill-down

From "share is dropping" to a SKU-level action.

BI shows trends but cannot decompose a gap by brand → segment → SKU → store, or distinguish KA from distributor logic.

meet troodon

An AI commercial analyst that scales the best of your team.

Troodon ingests your market and client data, applies the same diagnostic methodology your top analysts use, and returns a prioritized list of opportunities — for every customer, every category, every week.

What Troodon does

Diagnoses category × customer in two layers — market + client
Detects 9 gap types: share, price, distribution, coverage, stock, assortment, segment
Prioritizes opportunities by revenue impact
Adapts output to persona — Planning Lead, KAM, Field Rep
Differentiates KA accounts from distributors

How it's different

Methodology baked into the agent — not just Text-to-SQL
Versioned data dictionary + dynamic value discovery, robust to schema drift
SQL guards: SELECT-only AST validation, row limits, schema introspection
Knowledge Base reranked with Cohere v3.5 for analyst-grade answers
Modular by data maturity tier

Not in scope (yet)

No predictive forecasting — SageMaker is post-MVP
No write-back to CRM or SAP
No real-time ingestion — matches the monthly cadence of source data
No execution recommendations or autonomous decisions
Human in the loop is mandatory

architecture

Five layers — four shipped, one in development.

01	Application	Branded conversational UX with persistent history, thumbs-up/down feedback loop, and persona-aware output. Desktop in MVP.	CloudFront · S3 Cognito · API Gateway
02	Backend	Two specialized agents (Knowledge + Text-to-SQL) on AgentCore. sqlglot SELECT-only AST validation, dynamic value discovery.	ECS Fargate · FastAPI Bedrock AgentCore · DynamoDB
03	Knowledge Base	Versioned corpus governed in its own repo: methodology, definitions, view schemas (YAML), and use-case artifacts.	Bedrock Knowledge Bases Cohere Rerank v3.5 · S3 Vectors
04	Data foundation	Materialized views per source, versioned dictionary, role-based access. Engagement adapts to client maturity tier.	Redshift · Glue · Lambda SAP / Scentia connectors
05	Agent prompt	Planned 5th layer: source-of-truth for system prompts, behavioral rules (R1-R16) and persona templates — versioned independently of code.	Prompt registry · Eval harness Few-shot library · Guardrails

Built end-to-end on AWS — production-grade today.

GenAI & Knowledge

Amazon Bedrock
Claude Sonnet 4.6 via Converse API

Bedrock AgentCore
Managed agent runtime + memory

Bedrock Knowledge Bases
Cohere Rerank v3.5 + S3 Vectors

Titan Embeddings v2
1024-dim embeddings

Compute & API

ECS Fargate
FastAPI orchestrator + agents

API Gateway (REST)
JWT-authenticated entry point

Amazon Cognito
User auth + JWKS validation

CloudFront + S3
Frontend distribution

Data & Storage

Amazon Redshift
Materialized views, Data API

DynamoDB
Sessions, messages, feedback

AWS Glue / Lambda
Ingestion from SAP / Scintia

S3 (raw / curated)
Data lake foundation

Ops & Security

CloudWatch
Logs, metrics, alarms

Secrets Manager + SSM
Credentials & config

IAM + KMS
Least-privilege + encryption

Terraform + GitHub Actions
Multi-env CI/CD

approach & methodology

A four-phase delivery that proves value before scaling.

Each phase has explicit deliverables, exit criteria and a customer validation step. We do not move forward without sign-off.

01

Discovery & validation

2 weeks

Stakeholder interviews + workflow mapping
Data inventory + maturity assessment
Validation document signed by business team
MVP scope: 1 category × 2 client types

02

Data foundation

2–8 weeks*

Connectors: SAP, market data, distributor
Materialized views in Redshift, versioned dictionary
Access control + segment-level governance
*Depends on data maturity tier

03

Agent build & integration

4–6 weeks

Knowledge Base corpus + reranker tuning
Text-to-SQL agent with SQL guards (sqlglot)
Persona-aware orchestrator + feedback loop
Web app, auth, observability

04

Pilot, feedback & rollout

4–6 weeks

Pilot with 5–10 commercial users
Human-in-the-loop validation by domain experts
NPS + utility scoring · prompt iteration
Rollout plan + post-MVP roadmap

case study · in production

Molinos Río de la Plata.

INDUSTRY	CPG · Food manufacturing
PILOT SCOPE	Pastas Secas · Carrefour (KA) + Vemart (distributor) · 5 commercial users
STATUS	MVP delivered. Validation phase with MRP commercial team underway.
OUTCOME	A diagnostic that previously ran quarterly for top accounts now runs on every customer × category, on demand.

8 wk

SAP → Data Lake integration delivered

29.1%

Pastas Secas SPM volume covered (Carrefour)

9

Gap types automated end-to-end

18+

KB artifacts: context, schemas + use cases

7

Redshift materialized views in production

Sonnet 4.6

Anthropic model in production via Bedrock

industry benchmarks

GenAI in CPG & retail — what public research expects.

Independent research consistently estimates double-digit productivity gains and revenue impact from generative AI in commercial and category-management workflows.

McKinsey

10–15%

Sales productivity uplift. CPG identified among the highest-value-pool industries — ~3% of revenue.

BCG

30–50%

Category management and assortment cycle-time reductions in early GenAI deployments.

Gartner

> 70%

Of B2B sales orgs will adopt GenAI assistants for pipeline and account analytics by 2027.

AWS + Anthropic

< 3 s

Reference architectures show RAG + Text-to-SQL agents reaching production with sub-3s latency on Bedrock.

Numbers shown are public references for context. Troodon-specific impact is measured during the pilot phase against pre-agreed KPIs.

engagement model

Three packages, scaled to your data maturity.

We meet customers where they are. The same Troodon product is delivered through one of three packages, distinguished by the data foundation phase.

TIER 1	12–16 weeks	MOST COMMON	8–10 weeks	TIER 3	6–8 weeks
Data Lake from scratch. Customer without a consolidated data foundation.		Data Lake exists, needs alignment. Customer with a partial data foundation.		Mature Data Lake (Molinos profile). Customer with a production-grade data platform.	
Data Lake stand-up on AWS (S3 + Glue + Redshift) Connectors for SAP and 1 market-data source Data dictionary + governance baseline Full Troodon build & pilot Optional: dynamic-dictionary auto-discovery		Data assessment + materialized-view alignment Knowledge Base corpus build Full Troodon build & pilot Governance review + access-control mapping		Light data alignment + connector reuse Knowledge Base corpus + reranker tuning Full Troodon build & pilot Persona-specific output customization	

why dinocloud

The right partner to build Troodon with.

01

AWS Premier, with the right competencies.

Premier Tier Services Partner with AI/ML, Migration, Security, DevOps, MSP, CloudOps and Financial Services competencies.

02

GenAI in production, not in slides.

Troodon is live on Bedrock with Sonnet 4.6, Knowledge Bases, and SQL guards. We bring code, not concepts.

03

Methodology that compounds.

We invest in artifacts that survive the engagement — dictionaries, prompt libraries, persona templates, governance baselines.

04

LATAM-native delivery.

Local language, local commercial nuances, local data sources (Scentia, Logyt, Sig2k). Familiar with retail B2C regulatory realities.

05

Co-sell ready with AWS.

Aligned with the LATAM Partner Industry Sprint. Partner Solutions Finder listing, FTR-ready architecture, joint GTM motion in flight.

06

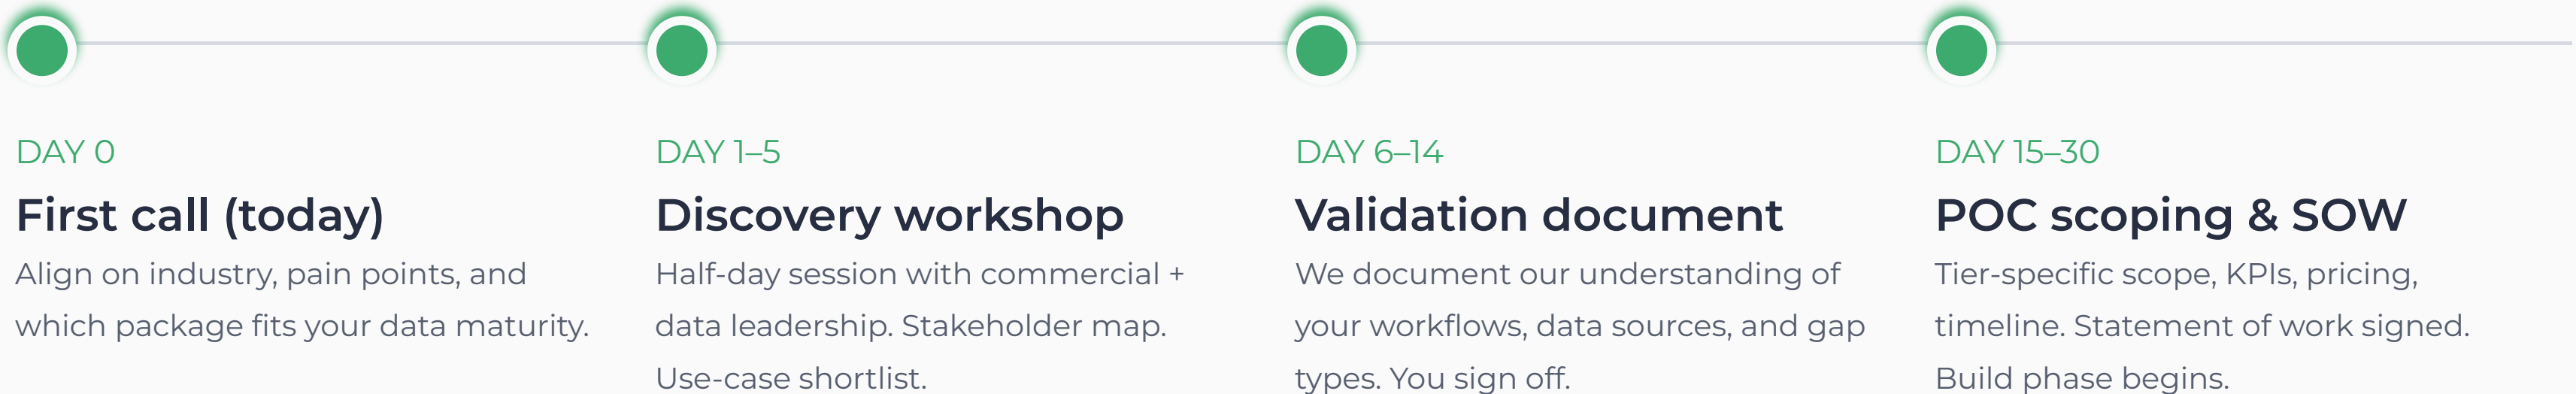
Human-in-the-loop philosophy.

Troodon augments analysts — it doesn't replace them. We design for trust, traceability and explainability.

next steps

The next 30 days.

A simple, predictable path from this first call to a signed Statement of Work.



Recommended ask. Schedule a **90-minute Discovery Workshop** in the next two weeks — no commitment beyond that session.

[Book the workshop](#)

Let's build it together.

Troodon turns your best analysts into a system. Every customer, every category, every week.

SPONSORS

Franco Salonia
Lead, Cloud & AI Solutions
franco@dinocloud.com

CO-LEAD

Lucas Borsatto
Commercial Lead, Troodon
DinoCloud

AWS

Premier Tier Services Partner
partners.amazonaws.com/dinocloud
Argentina · USA

WEB

dinocloud.com
Troodon · First Call Deck · May
2026
AWS LATAM Partner Industry
Sprint