

Weeks → Hours.

How Flipside deployed a full EVM blockchain data pipeline in under 2 hours using agent-driven automation — a repeatable playbook for every chain.

2hrs

FULL
DEPLOYMENT

20+

EVM CHAINS
DEPLOYED

42

AGENT
SKILLS

100+

DEFI
INTEGRATIONS

THE PROBLEM

EVM deployment was bottlenecked by complexity.

BEFORE · MANUAL

- 01 2–4 weeks per new chain deployment
- 02 Fragmented across siloed repositories
- 03 Manual AWS provisioning every time
- 04 Institutional knowledge in sparse docs
- 05 Every edge case solved from scratch

AFTER · AGENT-DRIVEN

- 01 Under 2 hours, end-to-end
- 02 One master skill orchestrates all steps
- 03 AWS + Snowflake provisioned automatically
- 04 42 skills encode institutional knowledge
- 05 Every edge case inherited on day one

THE AGENT LAYER

Three stages. One destination.

STAGE 01

2–4 weeks Manual

Engineers hand-built everything. Siloed repos, manual cloud provisioning, knowledge in people's heads.

STAGE 02

2–4 days fsc-evm Package

Standardized codebase cut redundant work — but deployment still required manual orchestration across dozens of steps.

STAGE 03 — NOW

1–2 hours Agent Skills

The full playbook encoded into structured skills. Agents orchestrate everything — engineers review and approve.

THE TOOLKIT

Years of engineering, encoded into three tools.

STREAMLINE

Per-chain infrastructure isolation

Each chain gets its own isolated AWS environment — Lambda, S3, API Gateway, IAM. Issues on one chain never cascade into another.

LIVEQUERY

Real-time node access from Snowflake

UDFs that call RPC nodes on demand directly from SQL — the agent's primary research tool for probing chain characteristics.

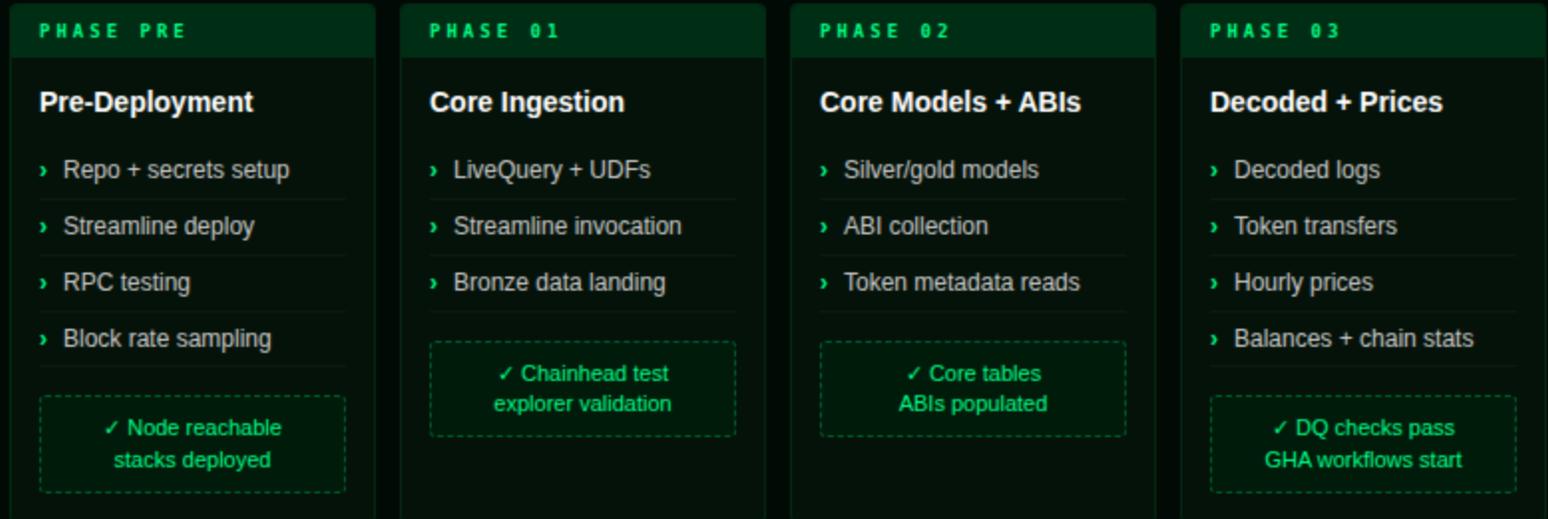
FSC - EVM

One codebase, every EVM chain

A single dbt package powering all chains through a 3-layer variable system. Node sampling auto-configures the pipeline. 70+ DeFi protocol integrations.

Four phases. Verification gates between every step.

INTAKE → chain name · node provider · known quirks



The compounding advantage of automation at blockchain scale.

DAY 1

Full coverage from day one

Manual builds took weeks for core schema alone. Agent deployments deliver the complete fsc-evm suite in a single pass.

AUDIT

Built-in verification trail

Phase gates catch issues early. Skills serve as living documentation — a clean audit trail of what was done, why, and how to proceed.

SCALE

Compounding returns

Each deployment surfaces edge cases that feed back into the playbook. Every new domain skill applies across every active chain.

RECOV

Graceful recovery

The agent handles real-world complications — infrastructure outages, provider quirks, partial state — picking up where it left off.

Pair-programming, not full autopilot.

THE AGENT HANDLES

- ✓ Snowflake resource provisioning
- ✓ AWS infrastructure stack deployment
- ✓ RPC node connectivity + chain sampling
- ✓ dbt model compilation, run, and testing
- ✓ Cross-validation vs. block explorer
- ✓ Status reporting at any checkpoint

THE ENGINEER REVIEWS

- › PR approvals at each phase
- › Phase transition authorization
- › Production Snowflake configuration
- › Node Vault secret verification

*"The engineer reviews and authorizes.
The agent executes and verifies."*

HyperEVM is the start. The system grows with every deployment.

Each new chain exercises the playbook further — without adding manual overhead.

NEXT →

Workflow Automation

Routine pre-deployment steps move into GitHub Actions so agents can focus on judgment calls, not mechanics.

NEXT →

Autonomous Phase Gates

Sleep commands replace manual waits. The agent verifies data landing and only escalates when a real decision is needed.

NEXT →

Expanding Domain Coverage

New skills covering DeFi primitives and deeper Snowflake/dbt expertise — each immediately applying across all active chains.