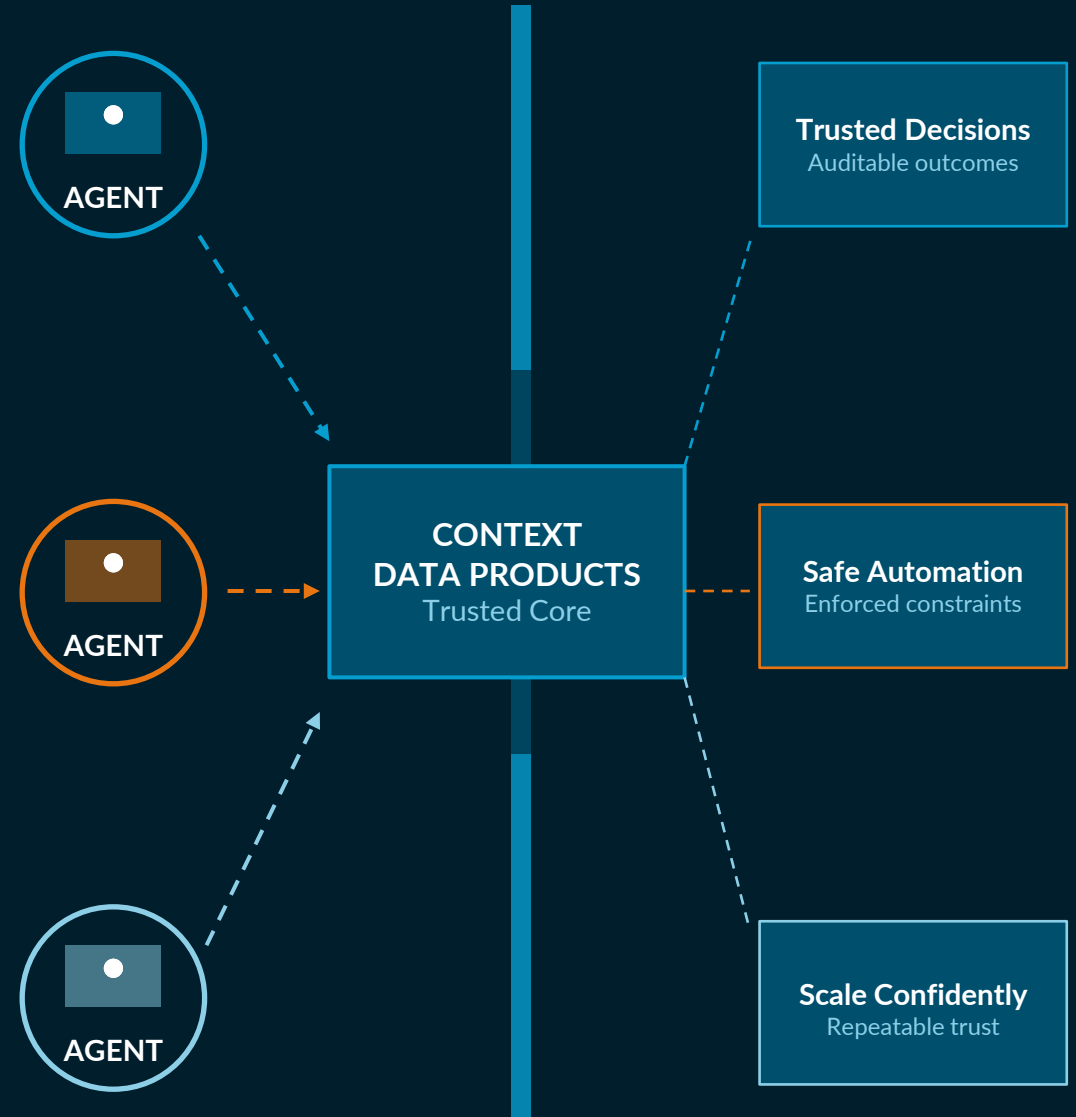


# Data for Agents: Building the Trusted Data Spine

Why Context Data Products are the foundation for agent-ready data

- AI agents amplify both value and risk
- Trust is no longer optional
- Context is the missing layer



# Presenters

---



---

**Maz Chaudhri**

**Practice Director,  
Intelligent Data & Automation**



---

**Jonathan Weiss**

**Lead Consultant,  
Intelligent Data & Automation**

# Setting the Stage

---



## Session Focus

---

- How agent-based AI changes what “good data” means
- Why Context Data Products are the unit of trust for AI
- How to scale them through a Trusted Data Spine



## Why It Matters

---

- AI agents amplify data ambiguity and risk
- Fragmented, ungoverned data breaks automation
- Trust, not models, is the bottleneck to AI at scale



## Your Takeaways

---

1. A clear definition of Context Data Products, and what makes them different
2. A practical mental model for building a Trusted Data Spine
3. Concrete starting points you can apply in your environment for your business



0 response submitted

## How is data primarily consumed in your organization today?

Scan the QR or use  
link to join



<https://forms.office.com/r/tVZURVJ7ph>

Copy link

Mostly through  
dashboards and  
reports

Via shared datasets  
or data lake access

Through APIs or  
reusable data  
products

Through agent or  
workflow driven  
automation

Treemap

Bar



1 of 4



# Agenda

---

## Transforming Data Foundations for AI Agents

- 01 What is Data Spine**  
Why Data Spine instead of Data Foundation
- 02 Analytics-Ready Data Spine**  
An overview of data foundation and governance capabilities
- 03 Why AI Agents Change the Rules**  
Limitations of analytics-ready data in agent-based decisions and actions
- 04 Context Data Product as a Unit of Trust**  
Evolution of data products to context data products
- 05 AI-Ready Data Spine**  
Context data products combined with trust capabilities
- 06 Approach to Get Started**  
Practical framework and startup plan

# Running Example: Proactive Customer Retention

A common business problem that exposes why agents need context, not just data.

## Business Situation

- Subscription-based business (SaaS / Media / FS)
- Customer churn is expensive and often detected too late
- Leadership wants AI agents to proactively identify risk and intervene

## Data Realty

- Analytics-Ready Data Products for: Usage, Billing, Support, Product exist
- They include traditional metadata (schema, owner, refresh, lineage, access)
- But business context lives outside the data product (dashboard, docs, tribal knowledge)

## Agent Goal

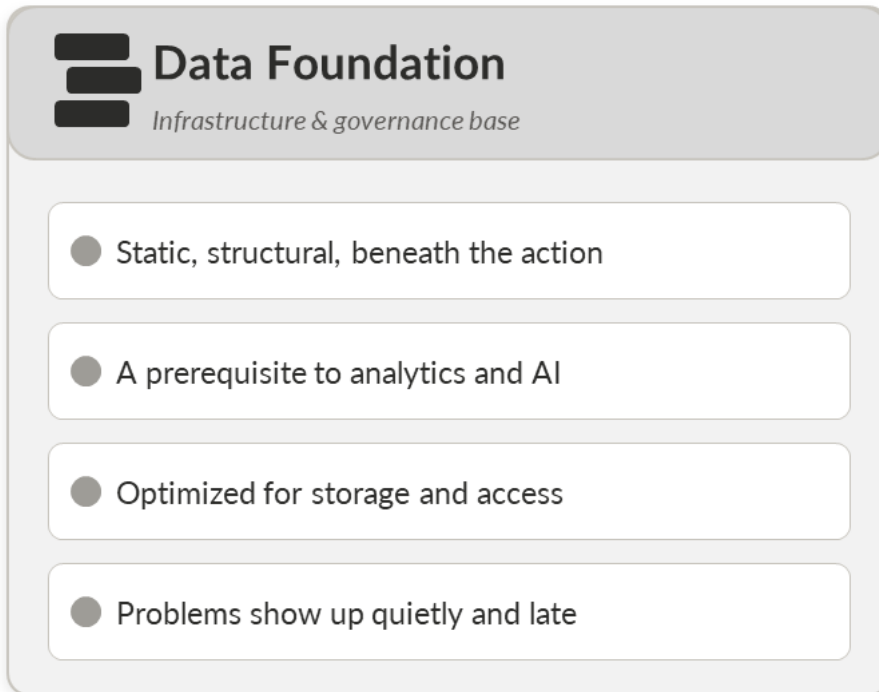
- Detect customers at risk of churn early
- Understand why they are at risk
- Recommend or trigger the right retention action: Offer, outreach, escalation, prioritization



# What is Data Spine

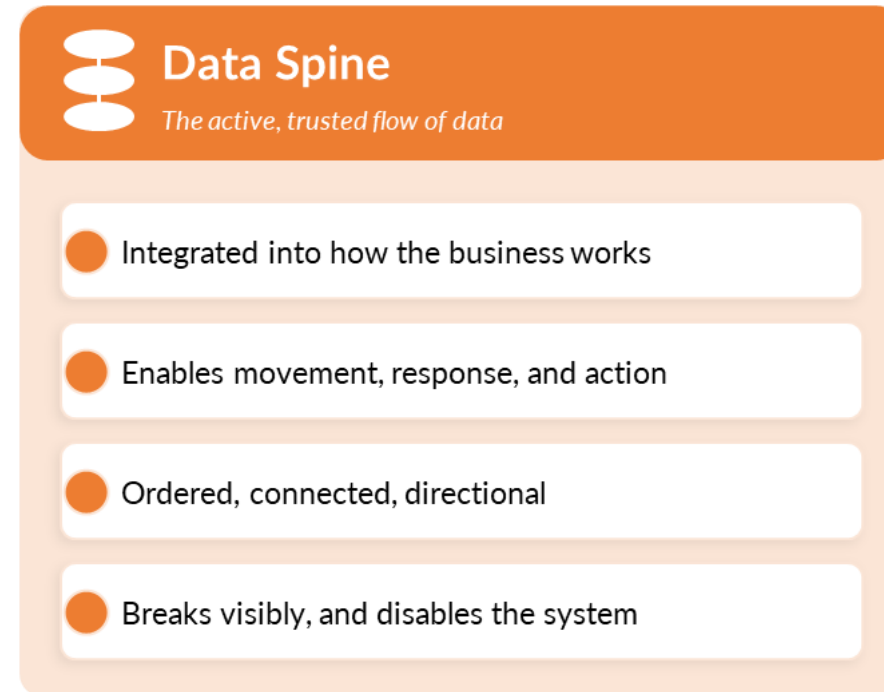
Why Data Spine – Not Just a Data Foundation.

A data foundation supports the business



- Stores customer, usage, billing, and support data
- Queried after issues appear
- Analytics first, action second

A data spine enables the business to move

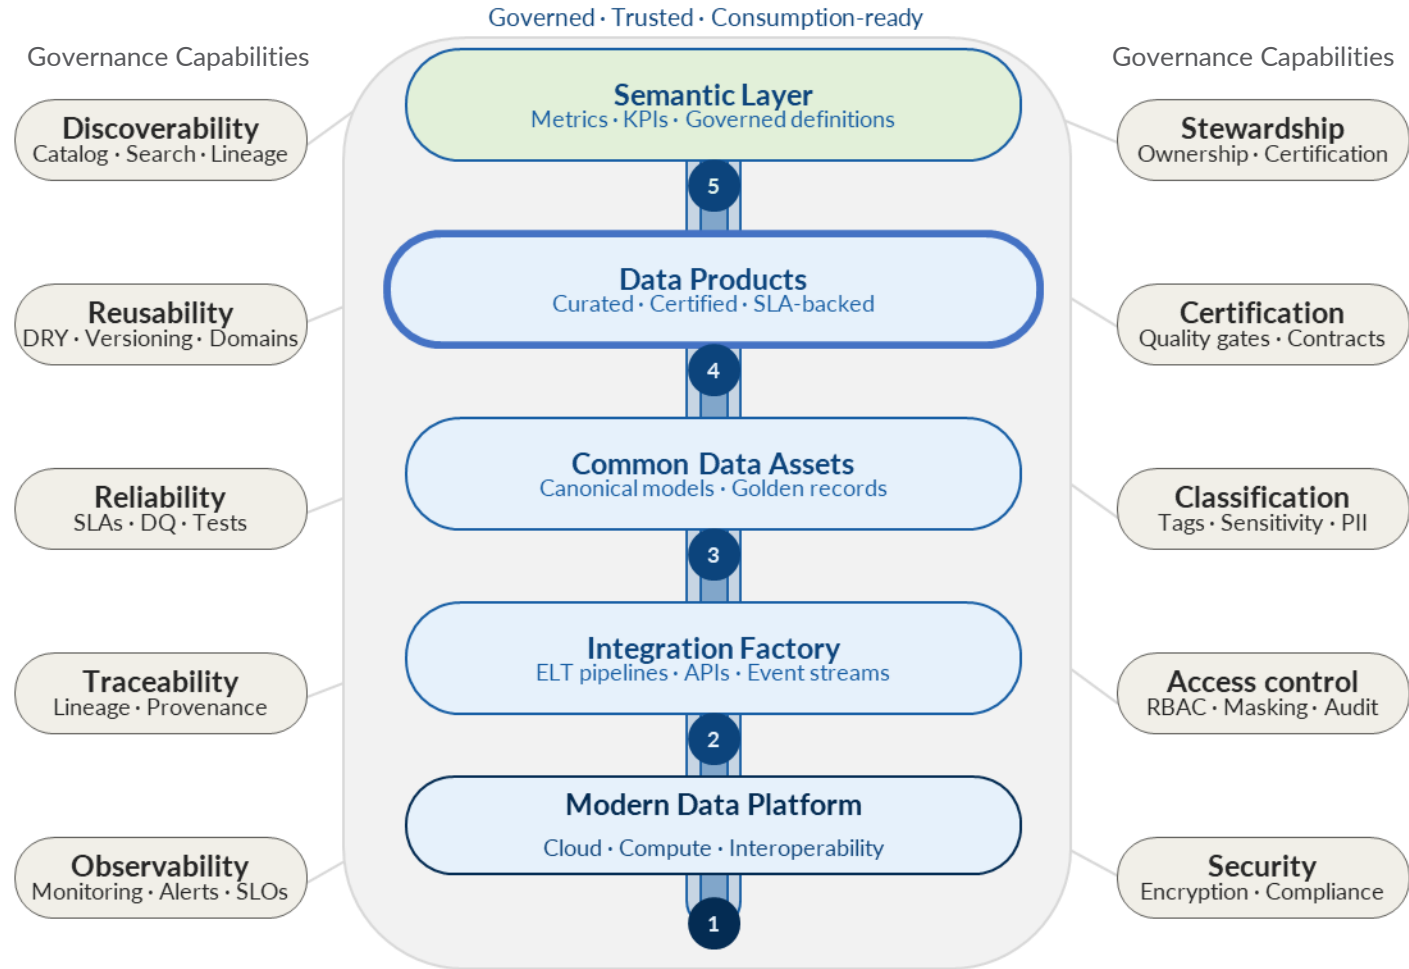


- Integrate real-time operational changes
- Feeds agents continuously
- Enables real time detection and response

# Analytics-Ready Data Spine

Data products sit at the center of the trusted data spine.

- 1 Scalable & elastic infrastructure
- 2 Adaptable integrations (CRM, billing, support)
- 3 Curated customer & subscription data
- 4 Consumption optimized data products
- 5 Centralized metrics/KPIs (churn rate, risk score)

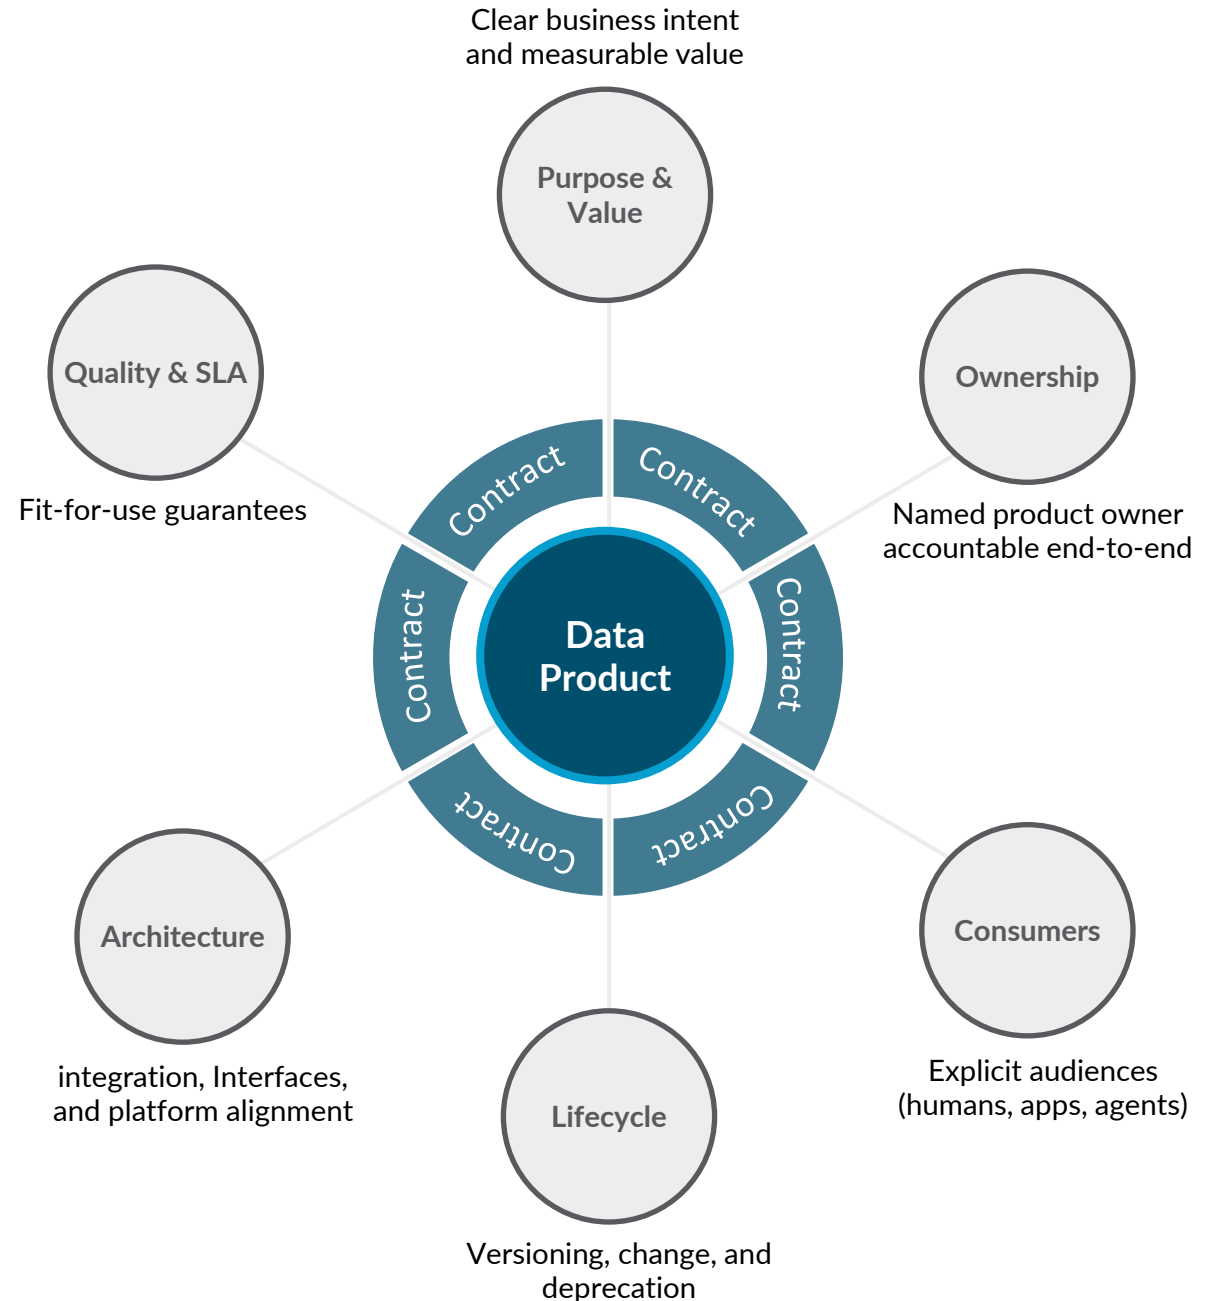


# Essential Components of a Good Data Product

A big-picture view of a data product.

- **Purpose:** Reduce churn
- **Owner:** Customer Analytics Lead
- **Consumers:** Analysts, ML models
- **Quality:** Freshness & accuracy SLAs
- **Lifecycle:** Versioned churn metrics
- **Architecture:** Secure access patterns

***Data contracts are the connective tissue.***

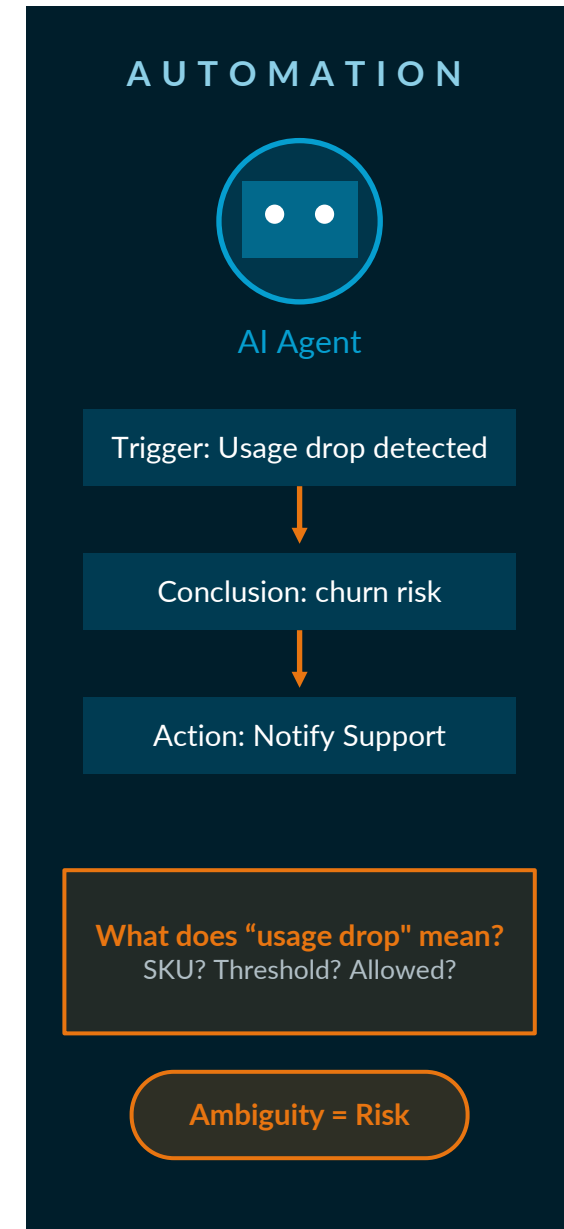
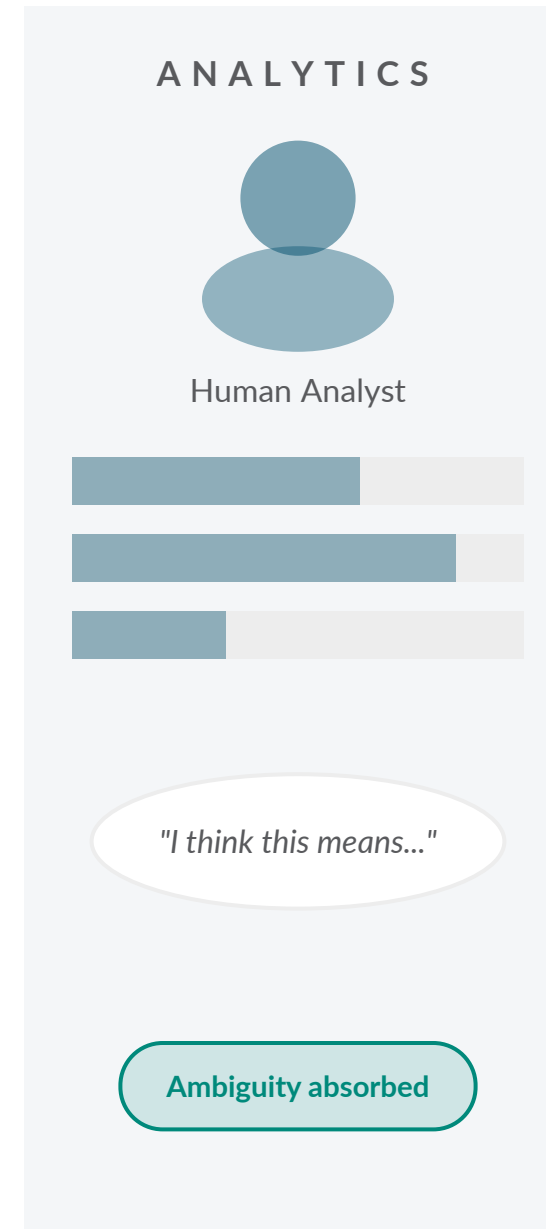


# "Good" Data Products Fail with Agents

Humans compensate for ambiguity. Agents can't.

- Analytics tolerates ambiguity
- **Agents require deterministic meaning**
- Trust, lineage, and policy become mandatory

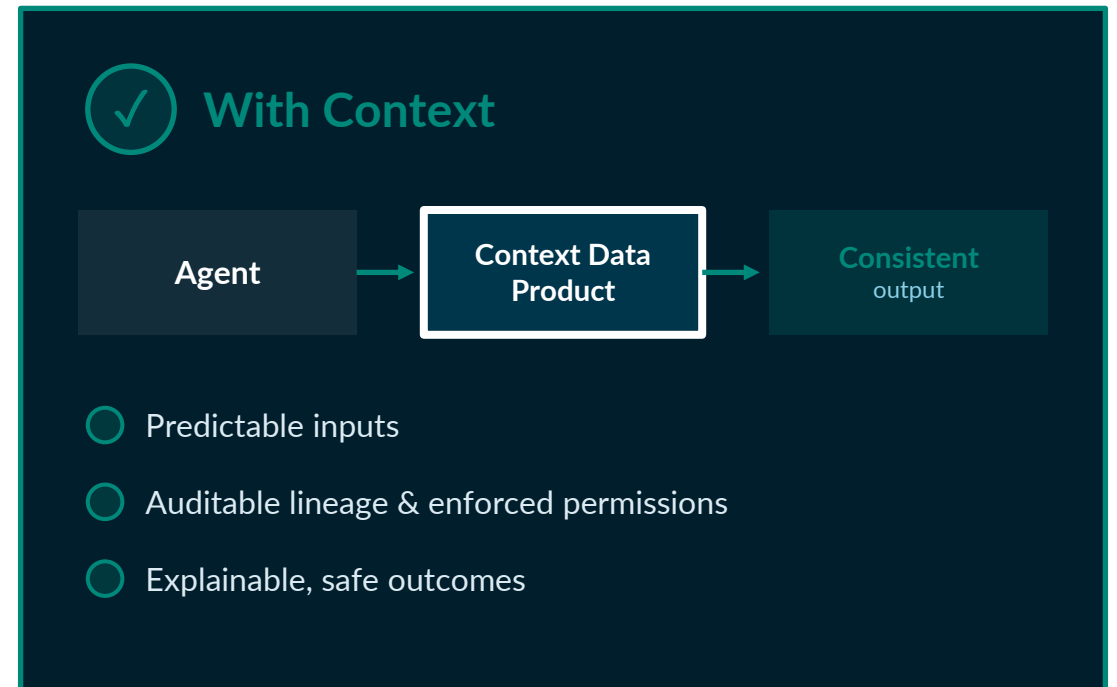
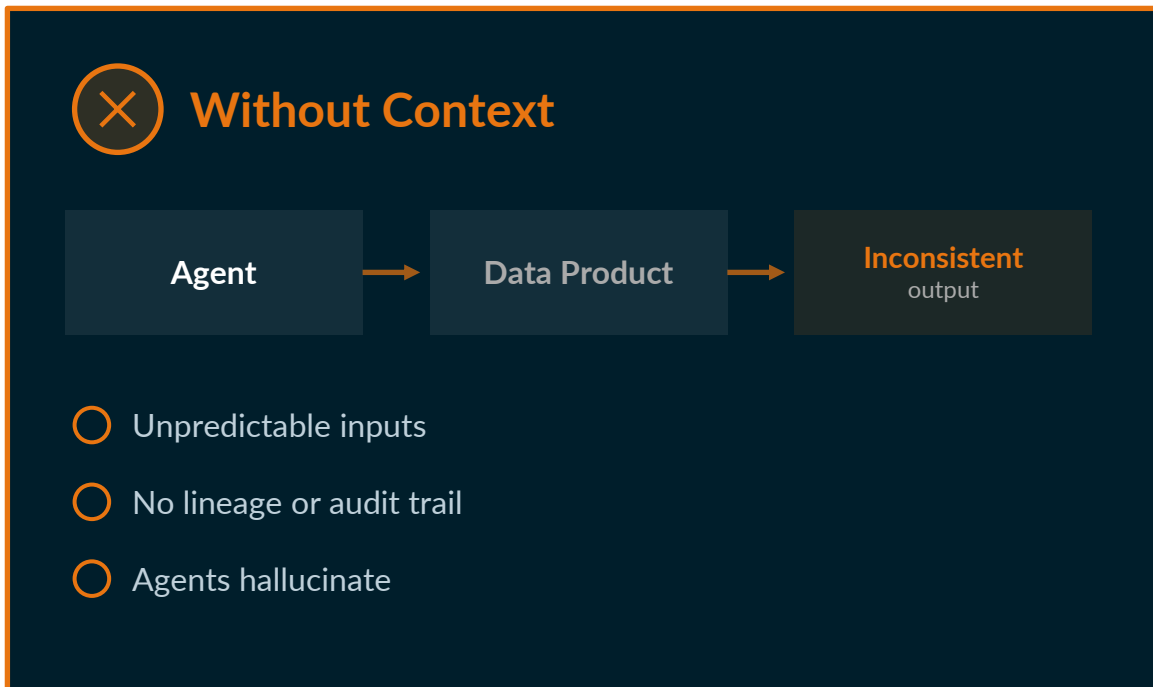
***Ambiguity becomes risk.***



# Agents Don't Just Need Data – They Need Context

Analytics-ready Products → human-readable context

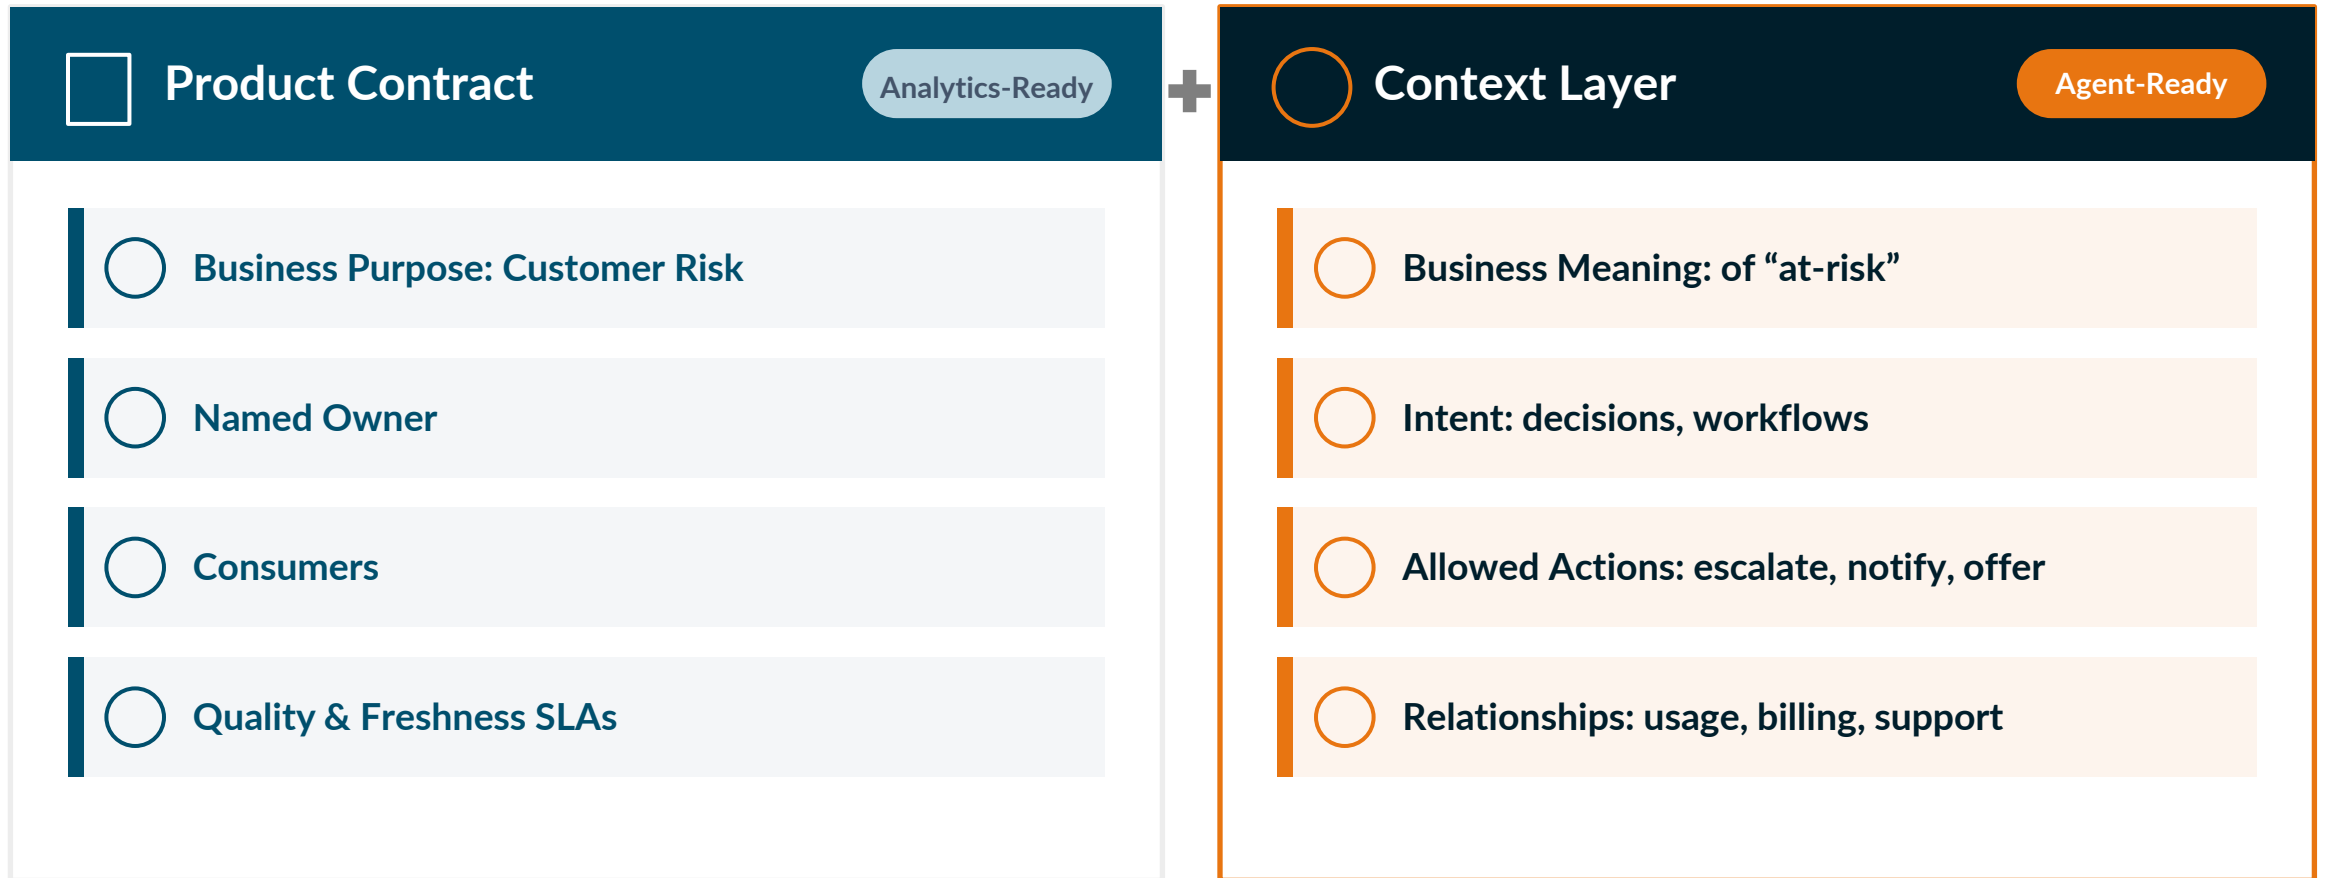
Context Data Products → machine-readable context for retrieve + summarize + recommend



Without context, agents hallucinate. **With context, agents behave.**

# Context Data Product is an Evolution of Data Product

Product Contract + Context Layer = Context Data Product



# Data Product vs. Context Data Product

Data Product ≠ Context Data Product

## Data Product

- Curated, reusable data asset
- Owner + SLA (technical focus)
- Schema & delivery focus (churn score table)
- Human interpretation required

Organizes data

## Context Data Product

- Curated Data + meaning + constraints
- Usage guidance + allowed actions
- Machine-readable and interpretable
- Agent-ready understanding

Organizes understanding

# What AI-Ready Data Actually Means

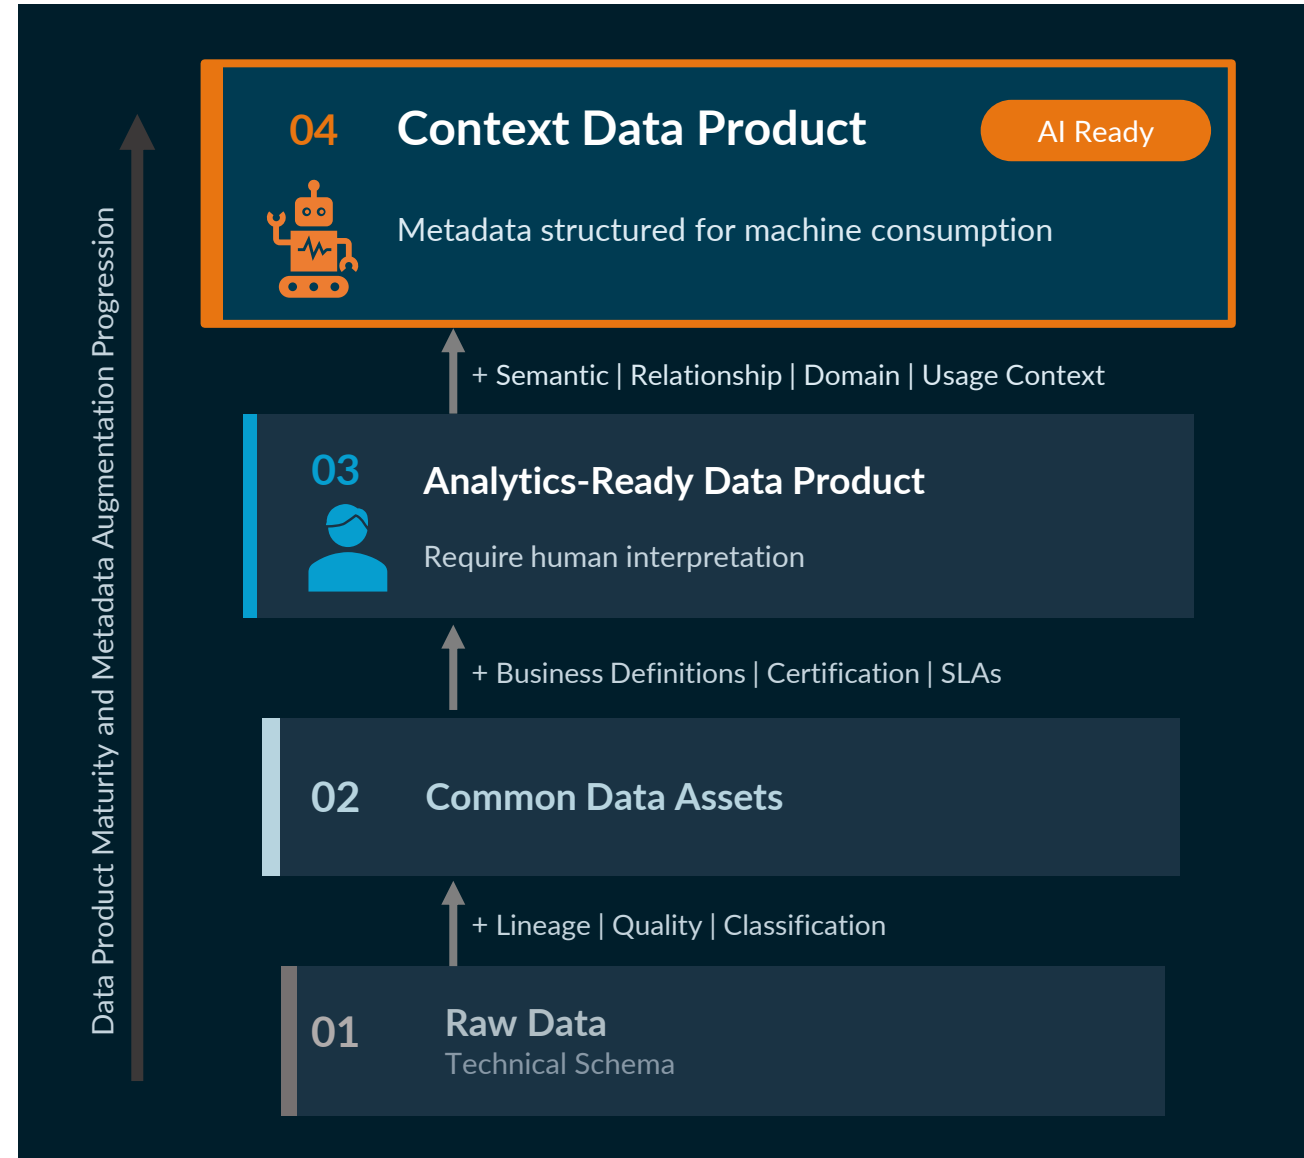
Clean and governed data isn't enough.

- Semantic meaning
- Relationship metadata
- Domain tagging
- Usage intent



- Traditional Metadata

**Context Data Product is AI-Ready Data**



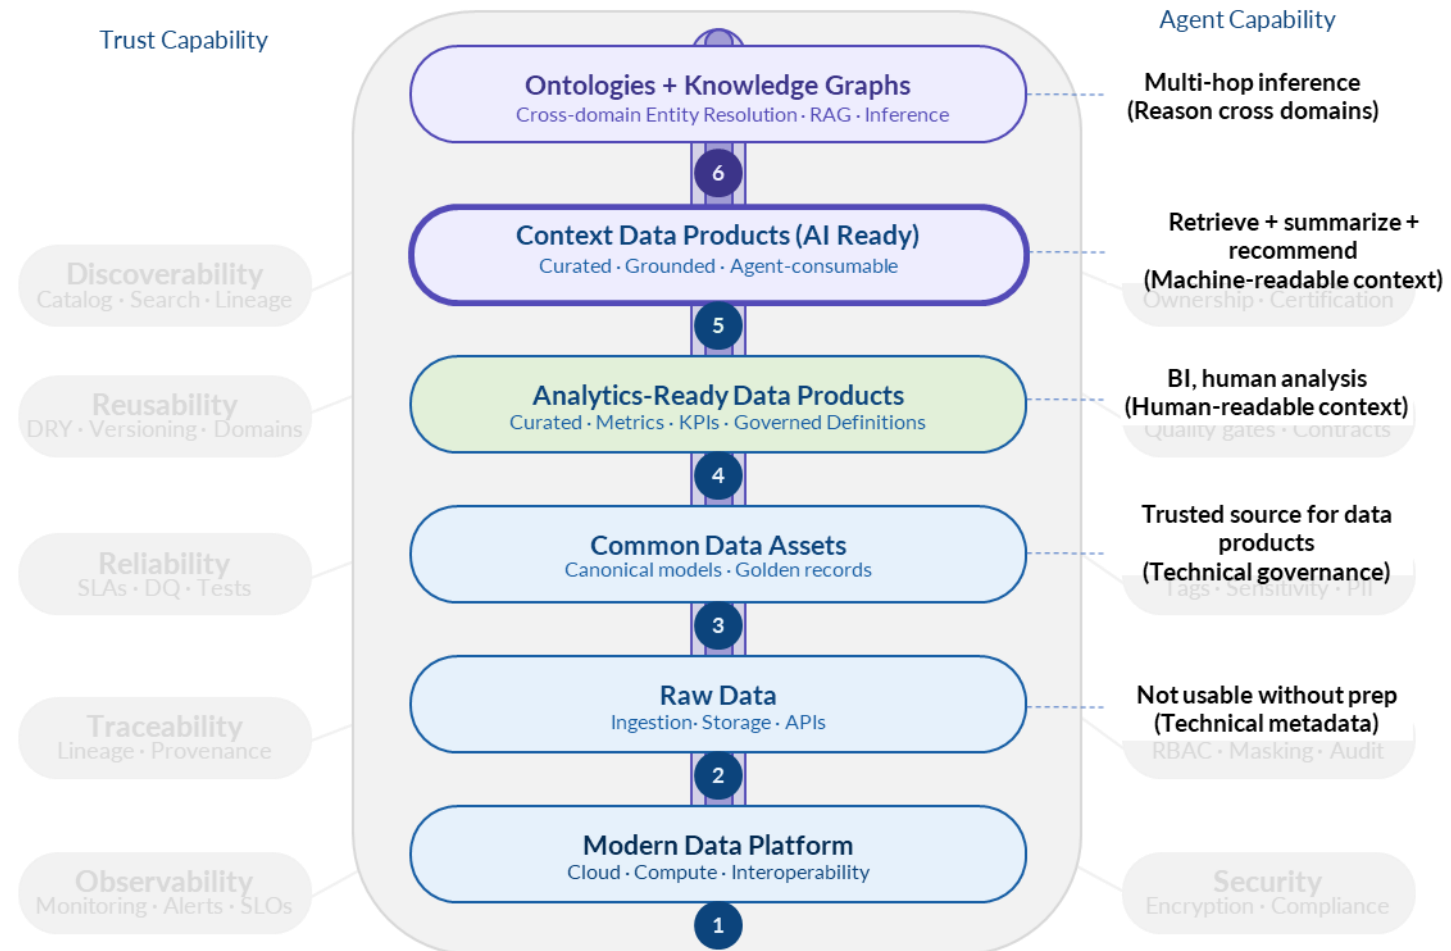
# AI-Ready Data Spine

## Context Data Products + Trust Capabilities

Ontologies and knowledge graphs are needed:

- *Cross-domain entity resolution becomes hard* (“customer” vs. “Account” vs “Household”).
- *You need multi-hop reasoning* (“customer risk due to supply constraints due to regional events”)
- *Semantic drift is frequent and costly* (“active user” or “revenue” varies across data products).
- *Agents must dynamically choose among many tools/data products* (tool planning)

**Context Data Products is the 80% Solution. Ontologies and graphs are for reasoning across many data products.**

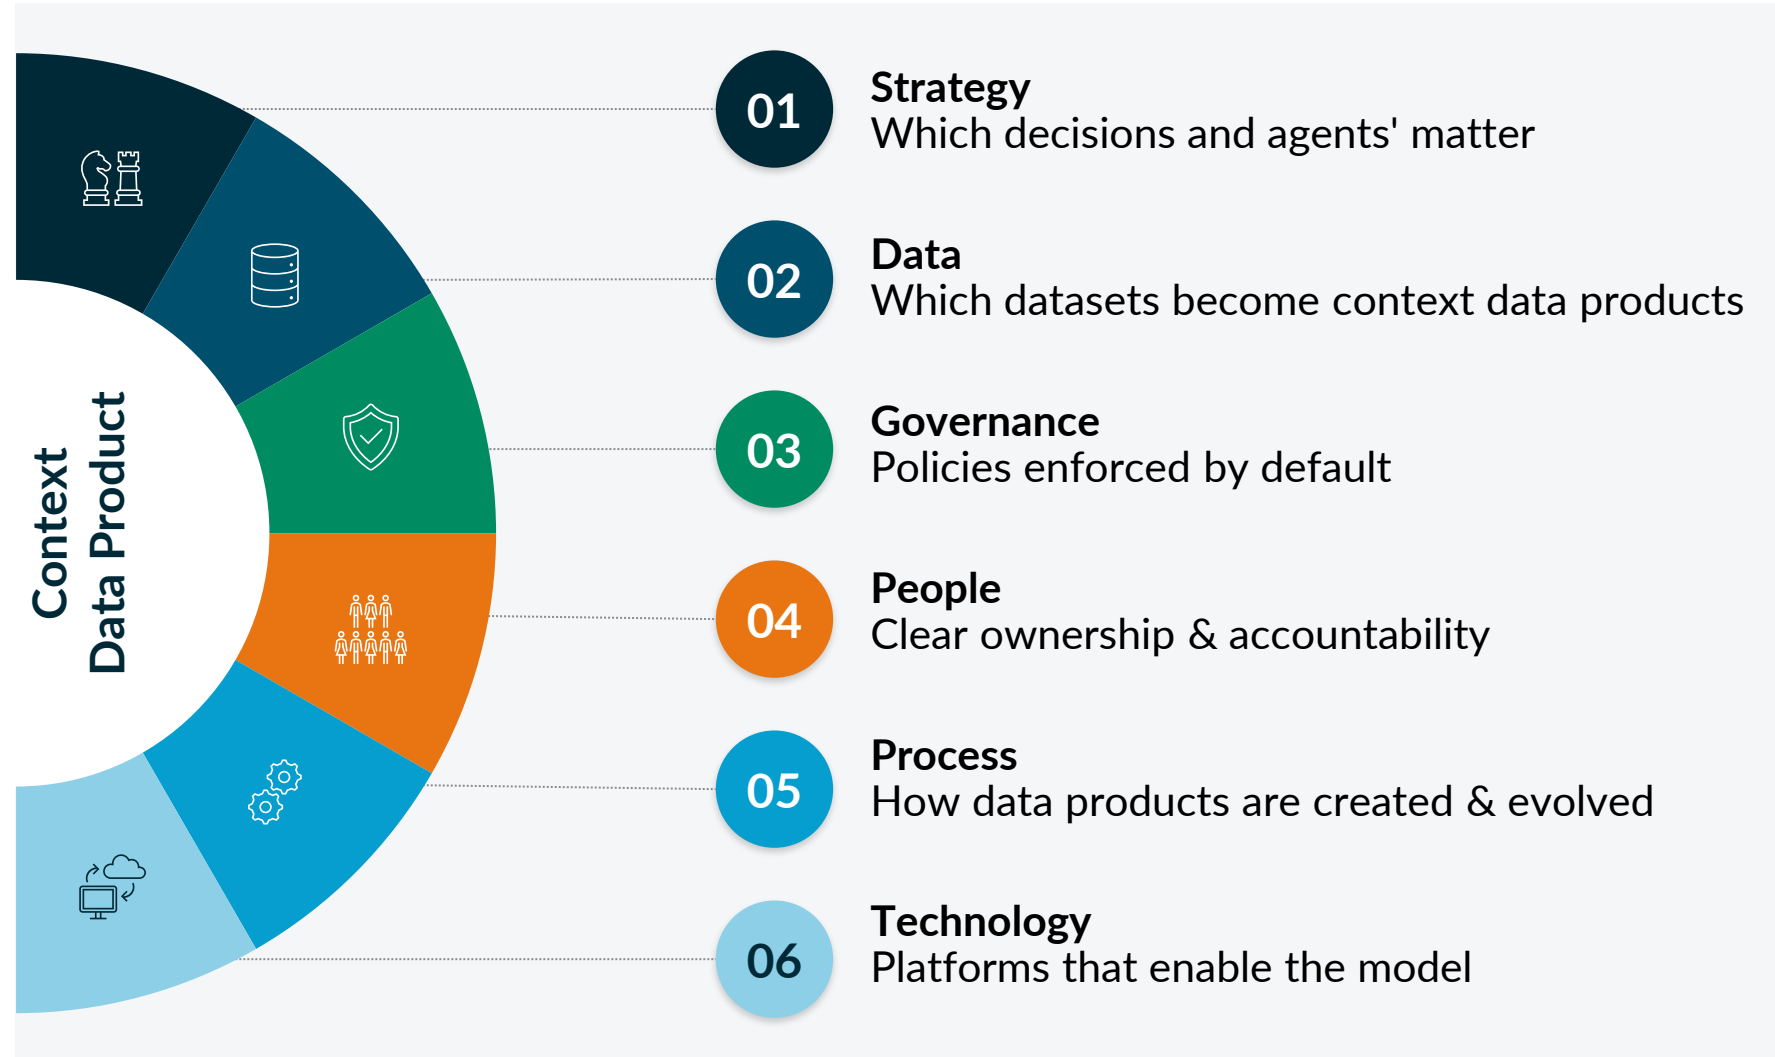


# Building the Trusted Data Spine

Framework to scale  
Context Data Products

**Align across six dimensions.**

It's an operating model lens  
that works no matter what  
your stack.



# First 90 Days: What Works

Start small. Build trust. Scale.

## ✓ Start Here

- 1 1-2 high-value agent use cases
- 2 1-3 context data products
- 3 Explicit owners & guardrails
- 4 Measure trust signals

## ✗ Avoid These

- ! Enterprise replatforming
- ! Tool-first programs
- ! Perfect model paralysis

Day 0

› Pick use case + name owners

Day 30

› First context data product live

Day 90

› Trust signals measured, spine growing

# AI Maturity Is a Data Trust Problem

No Trust → No Context → No Scale



Context Data Products are the bridge from AI experiments to real operations.

## FOUR ACTIONS TO START TODAY

- 1 Map where agents get data today
- 2 Promote one dataset to a context data product
- 3 Assign ownership, intent, and constraints
- 4 Build the spine iteratively