


23 June, 2026

From Data Chaos to Governed Intelligence: AI Agents in Action - Part 3

- Anita Ayyagari, Senior Success Guide
- Rashmi Prabhu, Senior Success Guide
- Shivam Saxena, Success Guide




- Today's Webinar is scheduled for **1 hour**
- The session will include a webcast and then your questions will be answered live at the end of the presentation
- All dial-in participants will be muted to enable the speakers to present without interruption
- Questions can be submitted to "All Panelists" via the **Q&A option** and we will respond at the end of the presentation
- The webinar is **being recorded** and will be available on our [Success Portal](#) - where you can download the **slide deck** for the presentation. The link to the recording will be emailed as well.
- Please take time to complete the **post-webinar survey** and provide your feedback and suggestions for upcoming topics.



Bootstrap trial and
POC Customers



Enriched
Customer
Onboarding
experience




Product
Learning Paths
and Weekly
Expert Sessions



Informatica
Concierge




Tailored training
and content
recommendations



Success Portal

<https://success.informatica.com>




Communities & Support

<https://network.informatica.com>



Documentation

<https://docs.informatica.com>



University

<https://www.informatica.com/in/services-and-training/informatica-university.html>

Safe Harbor

Disclaimer: The information being provided herein is for informational purposes only. The development, release and timing of any Informatica product, service or functionality described herein remain at the sole discretion of Informatica and should not be relied upon in making a purchasing decision. Statements made herein are based on information currently available, which is subject to change. Such statements should not be relied upon as a representation, warranty or commitment to deliver specific products, services or functionality in the future.

From Data Chaos to Governed Intelligence

AI Agents in Action - Series 3

Scaling Production-Grade GenAI with Metadata-Driven Trust and Model Context Protocol (MCP)

Where data
& AI come to **LIFE**

Agenda

01 The Reality of Enterprise AI Chaos : The Eighty Percent Problem

02 The Agent-to-Enterprise Architecture

03 Model Context Protocol: Connecting Agents

04 Live Demos — Document Intelligence, AI Readiness and Customer Intelligence

05 Q&A and strategic closing

The Data Iceberg: Governance is the Full Picture



**Bridging the Gap: Informatica CDGC + MCP brings context to the full spectrum.
(From Data Chaos to Governed Intelligence)**

The Reality of Enterprise AI Chaos

80%

OF AI VALUE IS DATA-BOUND

The Production Gap

Transitioning from experimental models to production-grade AI agents requires absolute data certainty. Without proper rails, autonomous agents operate in "chaos," executing decisions based on unverified, stale, or hallucinated facts.

The Path to Scale

To succeed at scale, enterprises must shift from manual data curation to automated, metadata-driven architectures that govern context in real-time.

What Is Agent-to-Enterprise (A2E)?

MCP Server is a standardized protocol that lets AI models securely connect to and interact with external tools, data sources, and services."

Agent to Enterprise is an Autonomous AI agents that take action across enterprise systems — from CRM to code — to complete complex, multi-step business tasks end to end."

Agent


LLM reasoning engine — interprets business questions, selects tools, synthesises answers

MCP

Model Context Protocol — standardised, secure way for agents to call enterprise systems

Agent + Protocol + Governance = A2E

MCP: Connecting AI Agents to Trusted Enterprise Data




MCP securely connects AI agents directly to Informatica's trusted, governed enterprise data in real-time.

MCP Enterprise Impact Matrix

Architectural Layer	Traditional Challenge (Without MCP)	Agentic Outcome (MCP-Enabled)
Metadata Search Layer	THE BLIND SPOT Agents reason over unclassified, stale documents with zero context.	DYNAMIC GROUNDING Agents instantly locate certified assets with real-time catalog context.
Data Quality Layer	HALLUCINATION RISK Agents process corrupt, stale, or incomplete data feeds.	TRUST SCORING Continuous data quality validation; agents bypass low-confidence sources.
Marketplace Layer	COMPLIANCE BOTTLENECK Manual security clearance takes weeks; unmonitored PII exposure.	AUTOMATED GUARDRAILS Self-service data contracts enforce dynamic masking and audit logs.
Customer ID Server (MDM)	IDENTITY CRISIS Fragmented systems create duplicate records and broken context.	GOLDEN RECORD RESOLUTION Universal entity resolution feeds agents a single customer truth instantly.

CDGC Metadata Catalog Search-MCP



settings_suggest Key Capabilities

storage

Metadata Retrieval

Returns crucial metadata including asset name, type, description, and unique ID for grounding.

psychology

Contextual Grounding

Provides AI models with necessary enterprise context to generate trustworthy and relevant insights.

auto_awesome

AI-Powered Classification

Leverages CDGC's automated classification (e.g., PII or financial) to associate business terms.

business_center Use Cases

gavel

Data Governance

Identify and report on tables containing sensitive data like PII.

anchor

LLM Grounding

Provide exact name and description of data assets before execution.

search

Data Discovery

Understand the full context of a data element across the catalog.

person_add

Onboarding

Summarize technical and business descriptions for new users.

Three People : One Governance Problem



Sarah Chen

Head of Claims Operations

BUSINESS OWNER

"Why does validating one claim take 3 days?"



Marcus Webb

Data Steward

GOVERNANCE

"I can't certify data I can't see."



Priya Sharma

ML Engineer

TECHNICAL

"I've been blocked for 3 weeks. Nobody will clear my datasets."



DEMO

AI AGENTS IN ACTION



Informatica

© Informatica. Proprietary and Confidential.

AI AGENTS IN ACTION • 2026

Case Study

Unstructured Data Governance • MCP Integration using Informatica• Curation Agent

The Use Case: Governing the 80% Nobody Can See

Unstructured documents are the data your AI agents are already reading — without governance context.

THE VOLUME PROBLEM

80% of enterprise data is unstructured

Policy PDFs. Claim forms. Compliance reports. Vendor contracts. Audit documents.

Not in databases. No schema. No column-level lineage. Invisible to your governance program.

THE TIMING PROBLEM

Your AI agents are already there

Claims adjuster AI reading policy PDFs. Contract review AI reading loan agreements.

These agents are in production — reasoning over documents with no governance context.

THE DEMO: WHAT WORKS TODAY

What CDGC + MCP enables now

Scan documents. Tag terms. Assign stewards. Certify.

Every query now retrieves governed context — definitions, steward identity, certification timestamp.

warning Current Limitation in CDGC

The process now is manual when we catalog PDF's/doc files we catalog but applying business terms, assigning stakeholders and classification of documents and tagging and certifying is now a manual process.


trending_up More Coming in This Space

Informatica is actively extending capabilities. What you see — term association, stewardship, certification — is the foundation. Deeper classification and lineage is in progress.

"Today's demo shows what is production-ready. We will call out the boundaries as we go."

What We Ship(Search Agent) — and What We Built(DataMint Agent)

The Architecture of Governed Context



Voice of the Stakeholder

- ❖ *GDPR audit **next week** - Which of our policy documents have been governed? Which ones contain PII?*
- ❖ *Claim's decision pending to validate coverage and produce an answer she can defend to a regulator.*



Marcus Webb
Data Steward • SureGuard

I have a GDPR audit next week."

Marcus opened CDGC. He searched. Three hours later he had a partial spreadsheet and an incomplete answer.



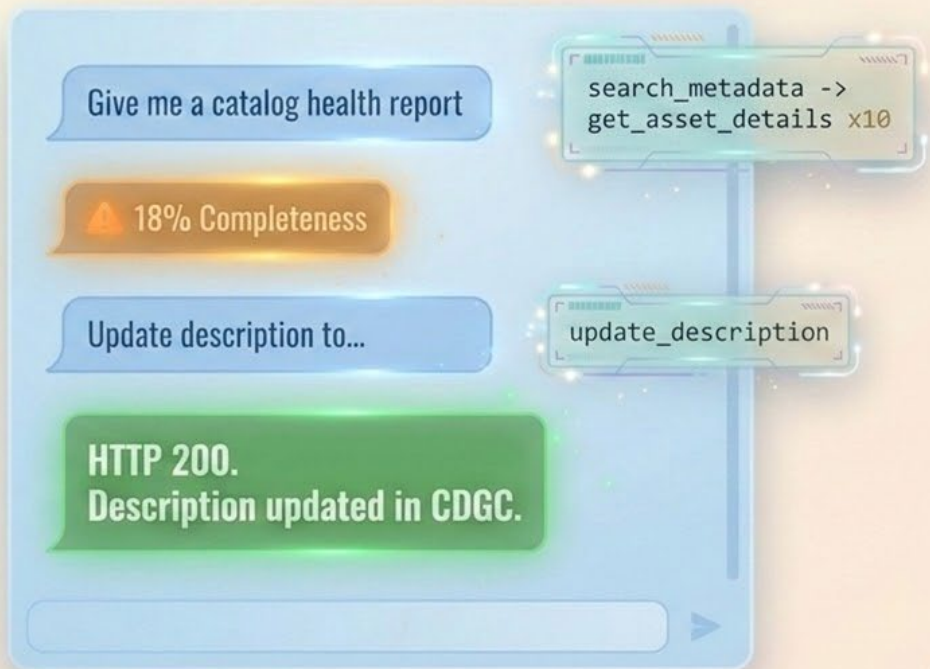
Sarah Chen
Head of Claims Operations

Claim #4521. Flood damage. Residential property.

Her AI agent needs to validate coverage and produce an answer she can defend to a regulator. Today that takes two to three days. A claims adjuster. Three PDFs. Manual cross-referencing.

Marcus's governance audit and Sarah's claim decision —
get solved in this session when we catalog and do minimal governance of our documents
[Using MCP without even logging into Informatica CDGC.](#)

Marcus: From 3 Hours to 30 Seconds



**The UI became the exception.
Certification remains the
only required human click.**

Sarah: The Defensible Claim

The Governed Agent

Coverage: YES.

Policy reference: Section 4.2,
policy_wording_life.pdf

Address: Validated (Cloud Address
Verifier)

M...ed 2026-06-15.

**Marcus Webb,
certified
2026-06-15.**

THE
AUDIT
TRAIL

The *Ungoverned* Agent

*The document likely covers... terms
may vary.*

**The agent is not the variable.
The governance is the variable.**

What Governance+MCP Delivered

STAKEHOLDER	BEFORE	AFTER
Sarah — Claims	2-3 days per complex claim	30 seconds, with audit trail
Marcus — Steward	Can't govern what he can't see	Documents catalogued, certified, traceable
CISO question	"I don't know"	Searchable, classified, policy-mapped
Regulator challenge	Empty audit trail	Full provenance to governed source
<i>Governance didn't slow Sarah/Marcus down. It's what made speed safe.</i>		

MCP + CDGC: The Governance Layer for Agentic AI

WITHOUT MCP

W

ar

hi

ro

st

ng

or

ou

vi

ti

to

ne

gg

le

_o

ff

Unverified Logic

AI Agent → Guesses → Unverified answer

Hallucination Risk

Compliance exposure via ungrounded outputs

Audit Failure

No trail → Impossible to govern or audit

WITH MCP + CDGC

AI Agent → MCP Server → CDGC

- **Business Glossary:** Standardized enterprise definitions
- **Data Classifications:** PII & sensitive tagging(Manual)
- **Governance Policies:** Real-time regulatory guardrails
- **Steward Info:** Direct links to data domain experts

✓ Trusted, traceable answer

"MCP is the nervous system. CDGC is the brain."

The Two-Week Audit Deadline: A Crisis of Visibility

“Is the data behind our SureGuard policy actually trustworthy—and can you prove it before the audit in 2 weeks?”



**Rashmi — Data Analyst,
Risk & Compliance**

- Needs to prove data meets the SureGuard policy before Q3 audit.
- Requires fast answers but lacks direct access to the catalog or data platform.



**Marcus — Data Steward,
Enterprise Data Office**

- Knows where everything lives but is constantly pulled into manual lookups instead of actual governance work.

Why the Process Breaks Down

Siloed Policies

Cannot link a policy to its data without manual navigation across multiple modules.

Invisible Governance Gaps

Certification status, missing descriptions, and PII gaps require entirely separate manual lookups.

Opaque Lineage

Upstream sources are displayed as meaningless system codes rather than business names.

Slow Access

Data consumers are blocked from remediation while waiting days for ticket approvals.

Human Bottlenecks

The catalog cannot answer complex questions without Marcus translating.

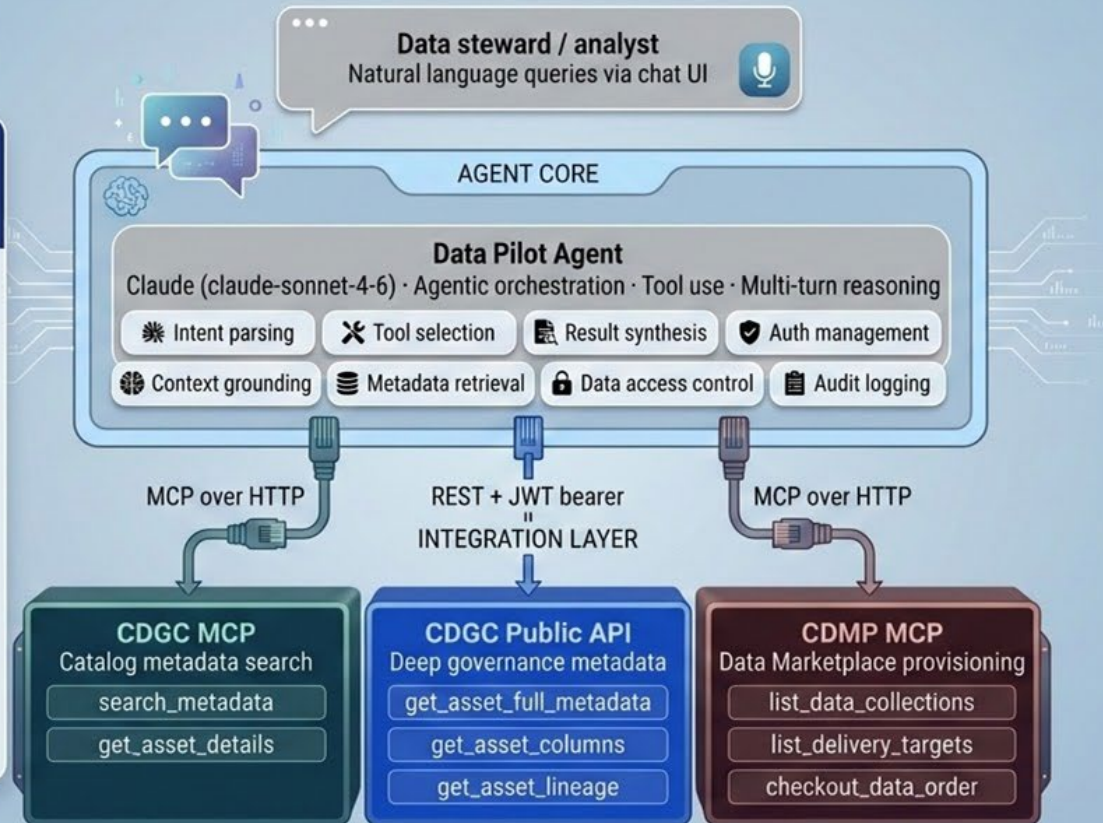
Manual Assembly

Audit preparation requires stitching together screenshots, emails, and spreadsheets.

Data Pilot Agent

Data Pilot Agent: Governed Intelligence

- Enables secure, policy-compliant data access.
- Orchestrates metadata discovery across systems.
- Automates governance checks (PII, certifications).
- Provides auditable lineage for agent actions.



The Transformation: Manual vs. Agentic

Before

Policy published, but unvalidated.

Gaps hidden until the audit.

Opaque, internal system codes.

3-5 day IT ticket wait time.

After DataPilot

Policy linked to real assets—verified in seconds.

15 specific gaps surfaced instantly across metadata and schema.

Clear upstream table with business-readable source names.

Requested and confirmed in the exact same conversation.



DEMO

AI AGENTS IN ACTION



Informatica

© Informatica. Proprietary and Confidential.