



27 Oct 2026


Bridging AI Assistants with Verified Master Records: Exploring the Customer Identification MCP Server for MDM SaaS

Ankit Tripathi, Success Architect

Housekeeping Tips




- Today's session is scheduled for **1 hour** and will feature a presentation followed by a Q&A session at the end.
- To ensure a smooth presentation without interruptions, all participants will be muted.
- Please submit your questions to "All Panelists" using the **Q&A panel**. We look forward to answering them during the live Q&A.
- This session is **being recorded**. The recording and slide deck will be available on our **Success Portal**, and a recorded link will be emailed to you after the event.
- We value your input! Please take a moment to complete the **post-webinar survey** to share your feedback and suggest future topics.



Bootstrap trial and
POC Customers



Enriched
Customer
Onboarding
experience



Product Learning
Paths and Weekly
Expert Sessions




Informatica
Concierge




Tailored training and
content
recommendations

More Information




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>

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.

Forward-looking statements



This presentation contains forward-looking statements about, among other things, trend analyses and statements regarding future events, anticipated growth and industry prospects, and our strategies, expectation or plans regarding product releases and enhancements. The achievement or success of the matters covered by such forward-looking statements involves risks, uncertainties and assumptions. If any such risks or uncertainties materialize or if any of the assumptions prove incorrect, results or outcomes could differ materially from those expressed or implied by these forward-looking statements. The risks and uncertainties referred to above include those factors discussed in Salesforce's reports filed from time to time with the Securities and Exchange Commission, including, but not limited to our ability to consummate the pending acquisition of Informatica on a timely basis or at all; our ability to meet the expectations of our customers; uncertainties regarding AI technologies and their integration into our product offerings; the effect of evolving domestic and foreign government regulations; regulatory developments and regulatory investigations involving us or affecting our industry; our ability to successfully introduce new services and product features, including related to AI and Agentforce; our ability to execute our business plans; the pace of change and innovation and our ability to compete in the markets in which we participate; and our ability to maintain and enhance our brands.

Agenda



- 01** Why AI agents need trusted customer data
- 02** MCP & the Informatica server catalog
- 03** Customer Identification + MDM-Provenance — the two servers

04 OAuth — secure agent access

05 Takeaways & Q&A

THE PROBLEM



Smart Model, Untrusted Customer Data

Ask an LLM “who is Bob Smith at Acme?” and it guesses — or reads one CRM and returns a fragment. The gap isn’t model intelligence; it’s trusted, resolved customer data.

Fragmented

The same customer lives in CRM, billing, support and legacy systems — each with a different key.

Duplicated & messy

“Bob”, “Robert J.”, “Bobby” — many spellings across systems, no single golden record.

Ungoverned

Agents acting on raw, unresolved data create duplicates and leak PII.

The real gap

Trust & identity resolution — not model intelligence. That’s what MDM delivers.

Agent + MCP + Governed Data = Trusted Answers

An AI agent is only useful if it can act on trusted data. Three parts make that safe:

Agent

The LLM reasoning engine — interprets the business question, picks the right tools, and synthesizes the answer.

MCP

The open, secure protocol that lets the agent call enterprise systems in real time — no custom integration.


Governed data

Informatica MDM SaaS — the single source of truth the agent is grounded in: golden customer records.

Autonomous agents act across your systems — grounded in MDM's trusted customer data, not guesses.

MCP: Connecting AI Agents to Trusted Customer Data

MCP securely connects the agent directly to Informatica's governed customer master data, in real time — no training data, no guessing.



No training data. No guessing. Every answer is resolved against governed master data — and provable to its source.

The Informatica MCP Server Catalog

20 shared (native) servers, fully managed by Informatica — plus composite servers you build.

CUSTOMER & DATA

★ Customer Identification (MDM SaaS)
Address Verification

GOVERNANCE & CATALOG

CDGC Metadata Search · Catalog Discovery · Catalog Enrichment · CDGC
Data Quality Score · Data Marketplace Order Mgmt · Data Provisioning

OPERATE IDMC

Job Mgmt · Log Mgmt · Audit Mgmt · IDMC Troubleshooting · Secure Agent
Mgmt · Connection Mgmt · IDMC Asset Promotion · IPU Mgmt · User Mgmt

COMPOSITE SERVERS — expose your own Application Integration processes as agent tools (★ our composite MDM-Provenance server).

Creating & Operating a Composite MCP Server

Creating and Operating Composite MCP Servers in Informatica IDMC

PHASE 1: APPLICATION INTEGRATION PREREQUISITES

Verify Process Eligibility in CAI



Before creating a Composite MCP server, ensure the target Application Integration is active, deployed to Cloud Server, and uses a REST/SOAP binding.

Enforce Standard Basic Authentication



The CAI process must be configured with Basic Authentication; anonymous processes or endpoints restricted strictly to the API Gateway are not supported.

**PHASE 2:
SERVER CREATION &
TOOL ANNOTATIONS**

Define Server Metadata & Actions



In AI Agent Engineering, define the server purpose, key capabilities, and target use cases, then import specific operations/actions from published Application Integration processes.

Annotate Tool Behaviors for AI Safety

-  **Read Only** Golsitic querios/reads without modifying the orristeranerd.
-  **Destructive** Performs irreversible updates (e.g., deleting tiles, dropping tables).
-  **Idempotent** Repeated calls with identiyal parameters produce the evact seme outcome.
-  **Open World** Connects to external, unpredictable third-party-web endpoints.
-  **Return Direct** Directly returns output to the end-user without LLM re-suminarization.

**PHASE 3:
POLICY & AUTHENTICATION
CONFIGURATION**

Select OAuth Grant Type

**Machine-to-Machine**
OAuth 2.0
(Client Credentials grant type)

**Human-in-the-Loop**
OAuth 2.1
(Authorization Code grant type with PKCE)

Enforce Operational Execution Limits

MCP servers automatically inherit default operational guardrails:



- Rate Limit:** Standard cap of 1,000 requests per minute.
- Payload Limit:** Maximum single payload size of 1 MB for request/response streams.

**PHASE 4:
SERVER INSTANCE
DEPLOYMENT**

Deploy Server Snapshot



Create Instance


Click "Create Instance" on the server overview page to snapshot current actions/policies and generate an active endpoint for AI client consumption.

Assign Isolated Context URL String

`https://.../mcp-servers/<id>/<context>`

Append a unique context string (up to 100 lowercase characters, numbers, and hyphens) to isolate endpoints by use case.

PHASE 5: END-TO-END RUNTIME ARCHITECTURE



Two Servers · Five Tools



Out-of-the-box
Customer Identification
identify & retrieve

search_master_record

get_master_record_details



Composite
MDM-Provenance
prove it

get_master_record_sources

get_master_field_provenance

get_source_record

Out-of-the-box answers “who is this?” · Composite answers “why should I trust it?”

Two ways in

Federate identity to your IdP, or register a client with Informatica's OAuth server.

Federated OAuth

Delegate to Okta / Azure AD; the agent presents a JWT token (requires SAML SSO).

OAuth 2.0

Machine-to-machine — client credentials, no human in the loop.

OAuth 2.1

Human-in-the-loop — Authorization Code + PKCE; Claude registers as a public client.

The agent never sees your IDMC credentials — it carries a scoped, expiring token you control, and you can revoke it.

Takeaways



Trust

Agents are only as trustworthy as the **data behind them** — MDM gives them one source of truth.

Grounded

Governed **person & organization golden records** — no hallucinated customers.

Provable

Provenance makes every field defensible — trace it to its source, with a trust score.

Start small

One server, one use case — dedup at the point of entry is the fastest win.



Thank
you