

Preparing Enterprise Data for **Oracle Agentic AI**

_____ with _____
Smart Data Platform



Table of Contents

Executive Summary	1
1. Oracle Agentic AI Changes the Data Requirement	2
1.1 From response generation to controlled action	2
1.2 The Current Oracle Agent Stack	4
1.3 The risk moves from answer quality to operational consequence	5
2. The Data Problem Behind Agent Failure	5
2.1 Four weaknesses in conventional pipeline	5
2.2 The agent feedback loop	6
2.3 Five faults that deserve priority	7
3. Define the Agent's Decision Context	8
3.1 The Six Elements of a Decision Context Framework	8
3.2 Six Control Domains That Make Agentic AI Trustworthy	10
4. A Five-Layer Architecture for Governed Oracle Agent Data	11
Layer 1: Source and change	12
Layer 2: Validation, integration, and mastering	12
Layer 3: Semantics, metadata, and policy	12
Layer 4: Oracle AI and agent services	13
Layer 5: Action and oversight	13
5. A Four-Step Path to Agent-Ready Data	15
5.1 Four-Step Method	15
Step 1: Define the decision and profile its evidence	16
Step 2: Connect, standardize, reconcile, and master	17
Step 3: Govern exceptions and approve trusted records	18
Step 4: Publish, trace, and monitor	19

6. Build Trust Before Expanding Autonomy	20
6.1 Minimum production controls	20
7. Measure Business Impact at the Workflow Level	21
8. Evidence from a Complex Enterprise Data Program	22
8.1 Project Scope	22
8.2 Business Situation	22
8.3 ChainSys Solution	23
9. Why ChainSys	24
9.1 ChainSys Smart Data Platform: A Unified Data Foundation	24
9.2 Decades of Data Management Expertise	25
9.3 Why Trusted Organizations Choose ChainSys	26
10. Conclusion: The Data Foundation Is the Agent Strategy	27
10.1 Start with one workflow	27

Executive Summary

Gartner forecasts that more than **40% of agentic artificial intelligence projects** will be cancelled by the end of 2027 because of rising costs, unclear business value, and inadequate risk controls. For Oracle programs, the first question is whether enterprise data is reliable enough for action.

Data weakness rarely appears as an obvious error. It surfaces when duplicate vendor records lead a procurement agent to the wrong supplier, an old customer status misleads a service agent, or delayed inventory data is treated as current. The agent may work as designed while its decision context fails.

Oracle is making these workflows practical. **AI Agent Studio** for Fusion Applications supports agentic applications. **Select AI Agent** handles planning, tools, reflection, and memory inside the Autonomous AI Database. **Oracle Cloud Infrastructure Generative AI Agents** provides retrieval, orchestration, guardrails, and human intervention. Newer database capabilities include Private Agent Factory, Unified Memory Core, Deep Data Security, and managed Model Context Protocol and Agent2Agent services.

ChainSys Smart Data Platform addresses gaps through discovery, profiling, cleansing, mastering, lineage, and governance. It gives Oracle agents a verified view of data, preserves traceability, supports human oversight, and helps enterprises scale automation with confidence. This whitepaper explains required architecture, controls, and operating practices.

Why It's Time to Strengthen the Data Foundation?

40%

of agentic AI projects are expected to be cancelled by the end of 2027. Gartner cites escalating costs, unclear business value, and inadequate risk controls.

(Gartner, June 2025)

45%

of business leaders identify data accuracy or bias as a leading barrier to scaling AI.

(IBM, January 2026)

8/10

companies cite data limitations as a roadblock to scaling agentic AI.

(McKinsey, April 2026)

1. Oracle Agentic AI Changes the Data Requirement


1.1 From response generation to controlled action

Traditional artificial intelligence finds patterns or predicts an outcome. Generative AI creates content. An agent works toward a goal and keeps going: it assesses the situation, develops a plan, finds the context it needs, uses approved tools, and checks the result. The work ends when the task is complete or a person needs to review it.

Predictive and generative models can both sit inside that workflow. A model might draft a recommendation. The agent can then check the supporting data, send an exception to the right owner, update a record, or start an approved transaction. Output quality is only part of the control problem. **The enterprise also has to decide what the agent can see, which tools are available to it, and exactly what it may change.**

Autonomy rises. Control evidence must rise with it.

Move forward only when the data, authority, and failure controls support the next level.
That makes autonomy a business-control decision, not merely a feature setting.



Agents need **current, complete, and governed information**. They also need clear permissions, business context, and traceability. Every action should be linked to the data, policy, and tool that supported it.

What an Agent Needs From Enterprise Data

Autonomy raises the bar. Six requirements, one thread tying them together.



Every action an agent takes should be linked to the data, the policy, and the tool that supported it.

Agentic AI does not mean giving agents unrestricted control. Data and control requirements depend on what each agent is allowed to do. A mature approach defines clear levels of autonomy:

Autonomy Level	What the Agent Does	Required Controls
Assist	Retrieves, summarizes, and proposes	Source citations, access controls, and human review
Recommend	Compares options and prepares a decision	Policy checks, confidence thresholds, and a named approver
Execute Within Bounds	Performs pre-authorized, reversible tasks	Tool allowlists, transaction limits, and audit logs
Coordinate	Orchestrates multiple systems and agents	Segregation of duties, workflow gates, and exception handling

1.2 The Current Oracle Agent Stack

Oracle's 2026 releases broaden where teams can build and run agents. The relevant capabilities now span applications, cloud services, and the database:



Oracle AI Agent Studio for Fusion Applications is used to build, connect, validate, deploy, govern, and observe agentic applications. It's current capabilities include an Agentic Applications Builder, workflow orchestration, content intelligence, contextual memory, and measurement of return on investment.



Select AI Agent is the agent framework within the Autonomous AI Database. It supports multi-turn work through planning, tool use, reflection, and memory. Available tools can include natural-language-to-SQL, retrieval-augmented generation, database procedures, and external REST services.



Oracle Cloud Infrastructure Generative AI Agents provides a managed environment for retrieval, SQL and retrieval-augmented generation tools, custom functions, orchestration, guardrails, retained context, and human intervention.



Oracle AI Database Private Agent Factory gives teams a no-code environment for database-centric agents. It can run in a container in the public cloud or on premises.



Unified Memory Core allows agents to reason across relational, vector, JSON, graph, text, spatial, and columnar data held in Oracle AI Database.



Deep Data Security enforces database access according to the individual end user. For private model execution, **Oracle Private AI Services Container** can also operate in disconnected environments where that is required.



Trusted Answer Search enables AI Vector Search to find a previously created report when the business needs an answer that has already been tested and can be repeated.



Autonomous AI Database Model Context Protocol Server makes approved Select AI Agent tools available to compatible clients. Through the Agent2Agent protocol, Oracle's managed **Agent2Agent Server** can expose teams of Select AI agents.



Select AI Supervisor Agent and agent teams make coordinated multi-agent workflows possible inside the Autonomous AI Database.

None of these products removes the need for data engineering or governance. In practice, they raise it. Even a secure, well-designed agent runtime can receive duplicate identities, vague business definitions, stale status fields, or data from a copy that was never approved.

1.3 The risk moves from answer quality to operational consequence

A dashboard puts the final interpretation in human hands. An agent may take the same value and use it to select a supplier, release a hold, create a service action, or set several downstream steps in motion. As the sequence grows, freshness and mastering matter more. So do permissions and the ability to trace what happened.

The working rule is straightforward. Raise the data threshold as autonomy increases. An advisory use case can tolerate some uncertainty when the sources are visible and a person still decides. Once the agent can transact, it needs stricter thresholds, narrower permissions, and a dependable stop mechanism.

2. The Data Problem Behind Agent Failure

2.1 Four weaknesses in conventional pipeline

Most enterprise pipelines were built for scheduled reporting, migration, or analytics in bulk. Agents place a different demand on them because the data may be used immediately in a sequence of decisions.

Latency creates state mismatch

A nightly load may miss a stock reservation, the closure of a service case, or the release of a credit hold. The agent therefore needs a freshness timestamp and an agreed service level. If the threshold is missed, there must be a safe fallback

Schemas do not carry business meaning

A field name alone cannot tell the agent which customer identity is authoritative. It also does not say whether an inventory balance is available-to-promise or which version of a policy is in force. Technical metadata needs a connection to the business definition, the owner, and the permitted use

Small defects compound

A duplicate supplier can point the agent to the wrong contract and payment terms, which then affects the purchase recommendation. In a workflow with several steps, each output can become the input to the next decision

Copies weaken control

Repeated extracts can scatter sensitive data and make lineage difficult to follow. The agent should work from an approved view or data product created for its purpose, not the copy that happens to be easiest to reach


2.2 The agent feedback loop

An agent feedback loop begins when unsupported output is written into an operational system. Later, another agent retrieves that output and treats it as evidence. In this paper, the phrase describes a system-design problem; it is not an Oracle product term.

Take a collections workflow. The agent reads an incomplete dispute record and concludes that the invoice is valid. That conclusion is written to the customer account. When a later agent finds the entry, it treats the entry as evidence and escalates collection activity. By then, the original data gap has been absorbed into the official history.

Agent feedback loop and control breakpoints

Unsupported output becomes more dangerous when it returns later as trusted evidence.



Break the loop before an unsupported conclusion becomes reusable enterprise evidence.

Five controls can interrupt this loop:

- Use only approved, versioned sources for retrieval.
- Before a tool is invoked, test the data for freshness, completeness, and confidence.
- Keep notes created by an agent separate from verified transaction evidence.
- When an action is uncertain, irreversible, or sensitive to policy, require approval.
- Watch for recurring failures, unexpected use of tools, and drift in outcomes.

2.3 Five faults that deserve priority

Duplicate identities. The same customer, supplier, product, or employee appears under several identifiers.

Incomplete or stale attributes.

Required fields may be missing, or a recorded status no longer reflects the actual business state.

No authoritative record.

Several applications hold conflicting values, but no approved survivorship rule determines which value should prevail.

Broken lineage.

A team cannot trace the recommendation back through its source, transformation, quality result, policy, and approval.

Unclear ownership.

Exceptions sit unresolved because accountability has not been assigned to a collector, steward, or approver.

The closer a defect gets to action, the more expensive it becomes. Finding it during a profiling run is relatively contained. Once a transaction has occurred, correction and investigation are needed, and notification may be required. If the same defect is reused in later decisions, both its reach and the recovery effort grow.

3. Define the Agent's Decision Context


Clean data is not automatically ready for an agent. A supplier record can be complete and still be unsafe to use. Its status may be out of date, the contractual relationships may be ambiguous, or the agent may lack permission to use a restricted attribute.

ChainSys recommends a **Decision Context Framework** for each agent workflow. This is the organizing framework used throughout the paper. It ties a specific business decision to the evidence and controls needed to support that decision.

3.1 The Six Elements of a Decision Context Framework

Decision Context Framework

Define the evidence and controls for one agent decision before expanding access or autonomy.



One decision



Defined evidence



Controlled authority



Traceable action

Decision purpose

Describe what the agent is allowed to assist with, recommend, or execute. Similar-looking decisions with a different risk profile should be explicitly left out.

Required evidence

Identify the entities, relationships, attributes, documents, and policy versions needed for the decision.

Quality and freshness thresholds

Set the mandatory fields and acceptable values, along with duplicate confidence bands, referential checks, and the maximum age of the data.

Authority boundary

Record the permitted data and tools. Add transaction limits and the conditions that send the work for approval.

Traceability requirement

Keep the source identifiers and transformation history, as well as data-quality results, retrieval evidence, tool calls, approvals, and the final action.

Failure behavior

Decide in advance whether the agent will hold, retry, use a safe default, request a correction, or escalate to a person.

The framework avoids a common mistake. Instead of giving an agent a broad enterprise data set, the team first agrees which records are fit for that particular decision.

3.2 Six Control Domains That Make Agentic AI Trustworthy

Governance for agentic AI must operate as an executable control system. It should not remain a policy document reviewed only by people. It must determine:

- Is the data approved for this purpose?
- Is the agent authorized to access and use it?
- Is the data current enough for the decision?
- Can the recommendation be traced to verified evidence?
- Does the action require human approval?
- Are Personally Identifiable Information (PII), purpose, and retention controls being enforced?


Control Domain	What It Requires
Metadata and Catalog	Business glossary, technical metadata, ownership, and data classification
Data Quality	Source-level profiling, validation, cleansing, and standardization
Mastering	One governed golden record for each customer, supplier, product, or employee
Lineage	End-to-end traceability across data transformations, retrievals, and tool calls
Security and Policy	Identity-based access, least privilege, PII masking, purpose restrictions, and retention rules
Observability and Audit	Monitoring of freshness, drift, and tool failures, with an immutable record of agent retrievals and actions

Meeting these requirements calls for an integrated data discipline, not isolated controls. ChainSys solutions bring discovery, profiling, cleansing, mastering, metadata, lineage, governance, and monitoring together across Oracle and non-Oracle environments. This creates a trusted and traceable foundation for agents to act within defined business, security, and human-approval boundaries.

4. A Five-Layer Architecture for Governed Oracle Agent Data

The architecture needs a clear separation between data preparation, agent reasoning, and transactional execution. Evidence should remain available from one end of that path to the other.

Five-layer architecture for governed Oracle agent data





Layer 1: Source and change


Systems of record remain where they are: Oracle applications, legacy enterprise resource planning systems, software-as-a-service platforms, databases, files, documents, and external services. ChainSys data integration capabilities connect to approved sources through supported application programming interfaces, database adapters, files, and other configured interfaces. The source capability and the freshness required by the workflow determine whether the connection uses batch, incremental, event, or change-data-capture patterns.

Layer 2: Validation, integration, and mastering

Incoming data passes through controlled stages. Mappings and transformations are applied through data integration and reconciliation capabilities of ChainSys, which also supports validation and reconciliation. Data quality and master data management capabilities examine critical fields, standardize values, identify duplicates, apply survivorship rules, and manage governed master records. It then applies the approved survivorship rules and manages the governed master records. Throughout the process, raw data should remain clearly distinguishable from data that has been validated and approved.

Layer 3: Semantics, metadata, and policy

Business terms, owners, classifications, lineage, quality signals, and dependencies are captured through Active Metadata Management. This layer also needs to register semantic definitions and document metadata. Where a vector index is used, its provenance belongs here too, alongside the policies that determine whether data is fit for the purpose of a particular agent.




Layer 4: Oracle AI and agent services

Oracle AI Agent Studio, Select AI Agent, OCI Generative AI Agents, and other approved Oracle services retrieve the governed context through controlled interfaces. The actual interface depends on the design. It could be a database view, an application programming interface, a retrieval index, a Select AI Agent tool, or a Model Context Protocol endpoint. ChainSys handles the preparation and governance of the data. Responsibility for the agent runtime, database enforcement, and application-level controls remains with Oracle.

Layer 5: Action and oversight


The agent reaches Fusion Applications or another transaction system through approved tools. Before a tool runs, a policy gateway should check the identity of the agent, the user context, authorization, and input validity. It also needs to confirm approval status, detect a duplicate execution, and enforce transaction limits. Irreversible or high-impact work should remain under human control until the evidence supports a wider level of autonomy.

The controlled path from source to action



Governance, security, data quality, and audit run across the entire path. None of them should be added as a final stage after the agent has already been built.

The Gateway Between Agent and Action



5. A Four-Step Path to Agent-Ready Data

5.1 Four-Step Method

This method starts with one business decision rather than a broad platform rollout. Each step leaves behind evidence that the team can review before giving the agent more autonomy.

1. Define and profile

Use case:

supplier recommendation

Capabilities: discovery, profiling, definitions, ownership

Business outcome: agreed scope and measurable data baseline

Gate: purpose, owner, evidence, and autonomy boundary approved

2. Connect and master

Use case:

duplicate supplier identities

Capabilities: integration, transformation, reconciliation, quality, MDM

Business outcome: trusted supplier record linked to its sources

Gate: mandatory attributes pass and mappings reconcile

3. Govern exceptions

Use case:

disputed customer data and restricted attributes

Capabilities: stewardship, approvals, metadata, lineage, access controls

Business outcome: unsafe data is held and approvals remain traceable

Gate: only approved records enter the governed data product

4. Publish and monitor

Use case:

inventory shortage and alternative fulfilment

Capabilities: governed publication, provenance, quality monitoring, exception response

Business outcome: agents receive current, approved, traceable context

Gate: access is limited and hold or escalation paths are tested

Step 1: Define the decision and profile its evidence

Practical use case:

Supplier recommendation for an urgent purchase.

The agent is allowed to compare approved suppliers and prepare its recommendation. Creating a purchase order or adding a new supplier remains outside its authority

How ChainSys supports the step

- The Data Integration solution connects the supplier, sourcing, contract, purchase-order, and performance sources that have been approved for the use case. For each one, the project team records the objects and fields, their relationships, the refresh pattern, and the source owner.
- The Data Quality solution profiles the evidence the decision depends on. That includes supplier identifiers and tax numbers, legal names, addresses, site status, payment terms, qualification status, and contract references. The tests cover completeness, uniqueness, valid values, date validity, and referential integrity.
- Through Active Metadata Management, the team registers business definitions, field owners, classifications, lineage, and intended agent use.
- A baseline is recorded before correction begins. It includes the duplicate rate, missing critical fields, stale-record rate, ownership that remains unresolved, and the current level of manual effort.

Business Outcome

A clearly defined data scope, assigned ownership, and measurable quality baseline. The business can see whether supplier data is reliable enough to support the agent's recommendation before development continues.

Readiness gate

Work moves forward only after the decision purpose and excluded actions are approved, along with the data scope, owner, and boundary of autonomy.

Step 2: Connect, standardize, reconcile, and master

Practical use case:

Two supplier records refer to the same legal entity.

One contains the legal name and current tax identifier. The other uses an abbreviated name and an address that is no longer current.

How ChainSys supports the step

- The Data Integration solution lands or streams approved source data into a controlled staging area. Mappings translate each source attribute into the canonical model, while parent-child relationships, source keys, and effective dates are retained.
- Data Transformation and Reconciliation capabilities normalize source values and compare source and target results. Reconciliation then compares source and target counts. Accepted and rejected records are counted separately, and material control totals are checked.
- The Master Data Management solution applies deterministic matching to strong identifiers such as tax or registration numbers. When names and addresses must be compared, configured probabilistic or fuzzy rules can be used. Confidence bands decide which matches can proceed automatically and which ones need a steward.
- The survivorship rules determine the source for each attribute of the golden record. Source authority may carry more weight than recency, for example, while completeness and approval status can also be considered. A record remains unmerged when the evidence is uncertain.

Business Outcome

A trusted supplier record connected to its original sources. Duplicate reviews are reduced, supplier identities are clearer, and the risk of selecting the wrong contract or payment terms is lowered.

Readiness gate

Mandatory attributes must pass validation and the mappings must reconcile. Every duplicate candidate above the review threshold also needs an owner and a recorded disposition.

Step 3: Govern exceptions and approve trusted records

Practical use case:

a service agent needs the customer's status, but the record contains a disputed address and a restricted identity attribute.

The agent is allowed to compare approved suppliers and prepare its recommendation. Creating a purchase order or adding a new supplier remains outside its authority

How ChainSys supports the step

- The Data Stewardship solution routes failed validations and ambiguous matches into role-based workflows. The task carries the rule that failed, the supporting source evidence, the fields affected, its severity, and the date by which it is due.
- A collector is able to correct the evidence. The steward can approve or reject the record, or return it for more work. For material changes, the history is retained in a form suitable for operational review and audit.
- Active Metadata Management provides classification, ownership, lineage, and policy metadata. Those signals help determine which attributes can appear in the governed view and which ones have to be masked or left out.
- Access across the ChainSys and Oracle layers relies on enterprise identity and least-privilege controls. Credentials for connected systems are retrieved from a managed vault at runtime and are not stored in the platform. Integration publishers authenticate through OAuth rather than stored usernames and passwords. Each of these controls should be mapped to the identity and authorization model of the Oracle service being deployed.

Business Outcome

Disputed, incomplete, or restricted data is prevented from reaching the agent. Exceptions reach the right owner, approvals remain traceable, and sensitive information is protected throughout the workflow.

Readiness gate

A record enters the governed data product for the agent only when it meets the requirements for quality, ownership, policy, and approval.

Step 4: Publish, trace, and monitor

Practical use case:

an inventory agent evaluates a shortage and prepares an alternative fulfilment option.

To do that, it needs available-to-promise inventory, current reservations, approved substitute items, the priority of each source, and the age of every balance.

How ChainSys supports the step

- ChainSys publishes the approved records through the governed interface chosen for the deployment. The interface might be a database view, an application programming interface, a batch feed, or a curated store. Whatever the pattern, it carries the business key and record status, the effective time and quality result, and the provenance required by the workflow.
- The Oracle agent retrieves that context through an approved Oracle tool or service. If the design uses Autonomous AI Database, Select AI Agent tools can be exposed through the managed Model Context Protocol Server. This is an Oracle runtime pattern. It is not a claim that ChainSys has a native ChainSys-to-MCP connector.
- Data Integration and Reconciliation solutions monitor data movement. Data Quality and Master Data Management solutions rerun applicable rules as records change. Active Metadata Management provided catalog, lineage, reporting, and governance visibility.
- A failed threshold leads to the response agreed during design. A stale balance, for instance, may put the recommendation on hold. Failed reconciliation can block publication, while a shift in data quality may create a task for a steward. An unusual tool or transaction pattern may be sent for security review.

Business Outcome

Oracle agents receive current, approved, and traceable data. Quality or reconciliation failures can stop an unsafe recommendation, while ongoing monitoring helps the business detect problems before they affect an operational action.

Readiness gate

The agent can see only approved context. Every material action can be traced, and the hold or escalation path for the workflow has been tested.

6. Build Trust Before Expanding Autonomy

6.1 Minimum production control

Every agent, service, and tool should be treated as its own digital actor. An agent does not need all the access held by the person who starts a request.

Begin with assistance. During that phase, test grounding and access, review the citations, and see whether the quality thresholds and failure behavior work as intended. User confidence matters as well. The agent should move into recommendation only when reviewers can explain why it selected an option. Execution within defined bounds comes later, and only for reversible, low-risk actions. Transaction limits, checks for duplicate actions, approval controls, and rollback behavior all need to be tested first.

With a coordinated team of agents, individual accuracy tells only part of the story. The handoffs and shared context need testing, along with delegation boundaries, ownership of tools, timeout behavior, and recovery when one agent fails. Select AI Supervisor Agent and Agent2Agent interfaces make these multi-agent patterns practical. They also create more places where context can change or authority can drift.

Minimum production controls

- Give each agent or runtime component a unique identity.
- Limit access to the data and tools required for its purpose.
- Retrieve only from approved sources and keep policies under version control.
- Check quality and freshness before a material action is taken.
- Send high-impact or uncertain decisions to a person for approval.
- Prevent duplicate transactions through idempotency or duplicate-action controls.
- Log the evidence, tool calls, approvals, and outcomes from end to end.
- Test what happens when the workflow is held, rolled back, or escalated.
- Monitor changes in data and policy, abnormal use of tools, and the resulting business impact.

7. Measure Business Impact at the Workflow Level

Measure	Baseline question	Intended direction
Decision reversal rate	How often does someone later correct a decision influenced by an agent?	Reduce
Exception resolution time	How long does a data exception remain open?	Reduce
Manual touch rate	Which cases still need a person, and why?	Reduce where policy permits
Unsafe-data interception	How many unfit records are stopped before the agent receives them?	Increase initially, then reduce upstream defects
Audit preparation time	How much time is needed to reconstruct one decision?	Reduce
New-source onboarding time	How long does it take before a governed source can be published?	Reduce
In-policy completion rate	How many eligible cases finish within the approved boundaries?	Increase safely

Each data improvement should lead back to an operational result. Better supplier mastering, for example, may cut duplicate reviews and reduce the chance of selecting the wrong contract. Stronger lineage often shortens an investigation. When mappings can be reused, adding a new source should take less effort. Continuous reconciliation has another practical benefit. It can find an upstream failure before the affected records reach the agent.

During the first weeks, the program should establish its baselines rather than promise final savings. Targets can be agreed once the team understands the workflow, data volume, exception rate, and control burden.

8. Evidence from a Complex Enterprise Data Program

From 40+ Applications to Trusted Enterprise Data: How a \$4 Billion Global Manufacturer Used the ChainSys Smart Data Platform

Note: Use it as a reference to creating trusted and governed Enterprise Data, not as a source for Oracle Agentic AI deployment.

A \$4 billion global engineering and manufacturing enterprise with more than 20,000 employees. The organization operates 25 facilities across six countries and serves data center, communications, and industrial markets.

8.1 Project Scope

The project established a central enterprise data management and data warehouse environment on a Hadoop Impala data lake. Its scope included:

- Integrating more than 40 legacy and cloud applications.
- Processing approximately 30 TB of enterprise data.
- Supporting historical, incremental, scheduled, on-demand, and real-time data flows.
- Cleansing and harmonizing data before loading.
- Maintaining global analytical reporting during the cloud transition.
- Preparing governed data for analytics and data science.


8.2 Business Situation

The client was moving from legacy systems and on-premises SAP ECC to Oracle Cloud Applications. Data was fragmented across 40-plus applications, five time zones, 1,000 data models, and 500 object sets. The business needed uninterrupted reporting and a scalable foundation for analytics and data science.

8.3 ChainSys Solution

- Established a layered enterprise data management and warehouse architecture.
- Connected legacy, SAP, Oracle Cloud, and other enterprise applications.
- Supported historical, incremental, scheduled, on-demand, and real-time data flows.
- Profiled and validated each data stream before ingestion.
- Cleansed, matched, merged, and harmonized master data.
- Reconciled data between source applications and the central environment.
- Applied workflow-based governance to data entering and leaving the platform.
- Enabled metadata management, lineage, cataloguing, security, and global reporting.

Illustration



9. Why ChainSys

Most platforms address only a portion of the agent-data problem. An integration tool can move the data, while a quality product finds defects and a catalog documents the assets. For an agent to be ready, those disciplines have to operate as one connected cycle. This is exactly what ChainSys Smart Data Platform offers.

9.1 ChainSys Smart Data Platform: A Unified Data Foundation

The ChainSys Smart Data Platform brings together the capabilities needed to prepare, govern, and deliver enterprise data for Oracle agent workflows. It connects source systems, transforms and reconciles data, improves data quality, creates governed master records, and maintains the metadata and lineage needed to understand how information was produced.

This creates a controlled path from operational data to agent use. Before an Oracle agent receives access to production data or approved tools, the implementation team should:

- Confirm the source systems, data formats, transformation rules, and reconciliation requirements.
- Review data-quality thresholds, matching logic, survivorship rules, and exception workflows.
- Define business terms, data ownership, lineage requirements, and governance policies.
- Validate deployment, data residency, encryption, secrets management, and identity integration.
- Decide how governed data and approved tools will be made available to Oracle AI Agent Studio, Select AI Agent, Oracle Cloud Infrastructure Generative AI Agents, and Fusion workflows.
- Set transaction limits, approval gates, and escalation rules for each use case.
- Assign responsibility for monitoring, service levels, incident response, audit evidence, and ongoing support.

These questions are not weaknesses to conceal. They are design decisions that determine what information an agent can use, which actions it may take, and when a person must intervene. They should be settled before production access is granted. Addressing them early gives the organization a clearer route to reliable business outcomes while maintaining control over enterprise data and agent activity.

9.2 Decades of Data Management Expertise



For over two decades, ChainSys has been at the forefront of data management innovation, providing proven solutions to the world's most complex data challenges. Our team of data experts brings an in-depth understanding of the technical, regulatory, and business complexities of managing data at scale.

Why Our Expertise Matters:

End-to-End Data Solutions: From data migration and integration to data governance and analytics, we cover all aspects of your data journey.

Industry Experience: We've successfully worked with a wide range of industries, including finance, healthcare, retail, manufacturing, and energy.

Proven Methodologies: Our best-in-class methodologies ensure that your data strategy is both robust and adaptable to the fast-evolving digital landscape.

Pro Tip:

Leveraging ChainSys's expertise gives you the advantage of working with seasoned professionals who understand the nuances of data management. Our consultative approach ensures that the solutions we recommend align with your specific goals.

9.3 Why Trusted Organizations Choose ChainSys:

Proven Track Record at Scale

From Fortune 500s to government agencies, ChainSys solutions have powered over 300+ successful enterprise implementations worldwide. Our platforms manage billions of records across complex ecosystems, ERP, CRM, cloud, legacy, and more, demonstrating scalability, speed, and success in high-stakes environments.

End-to-End Data Intelligence, Not Just a Catalog

ChainSys dataZen goes beyond traditional metadata management by offering a complete data intelligence suite, including data profiling, lineage tracing, quality scoring, access controls, and usage analytics all in one unified platform. That means fewer tools, faster implementation, and deeper insights.

AI-Powered Automation

Manual tagging and governance are a thing of the past. ChainSys brings AI and ML-driven automation to metadata discovery, classification, sensitivity tagging, and relationship mapping. This enables faster onboarding, better compliance, and smarter decision-making.

Enterprise-Grade Governance & Security

With robust role-based access, customizable approval workflows, and end-to-end lineage tracking, ChainSys helps organizations stay audit-ready and compliant with evolving global regulations—GDPR, HIPAA, CCPA, and more.

Our Differentiators

70%

Cost & Time
Reduction in Data
Ingestion Migration

99.9%

Data
Quality
Improvement

360°

Perspective with
Automated
Governance
Protocols

Low / No code

Approach for ease of
Development
& Rapid Deployment

10. Conclusion: The Data Foundation Is the Agent Strategy

10.1 Start with one workflow

Oracle's agent stack brings enterprise AI closer to live decisions and operational action. As that happens, the standard for data readiness changes. A record is not ready simply because it migrated or loaded successfully. The enterprise needs to know that it is authoritative and current, that the agent is permitted to use it for this purpose, and that the record can be traced to its source.

Leaders still need to ask, "Which agents should we deploy?" A second question is more immediate: "Which decisions can our data safely support today?"

ChainSys helps an organization answer that question through integration and data quality, master data, metadata, lineage, governance, reconciliation, and ongoing monitoring. Oracle provides the agent and database capabilities. ChainSys strengthens the decision context on which those capabilities rely.

Start with one workflow

Choose one Oracle agent use case where the business value is visible and the risk can be managed. Define a Decision Context Framework for it, then profile the underlying data. This reveals which controls are already in place and where a gap could affect a live decision.

Give your next Oracle agent evidence it can trust, controls the business can defend, and a clear route from a recommendation to responsible action.

Book a Demo

to work through one high-value Oracle agent workflow and see how ChainSys connects, masters, governs, publishes, and monitors the enterprise data behind it.



World's
Leading Brand
Trust
Our Data Solutions

Authors



Amarpal Nanda

President of EDM
amarpal.nanda@chainsys.com



Iswarya Tamilmani

Lead Technical SME . Marketing
iswarya.tamilmani@chainsys.ai



“

I am honored to serve as ChainSys's CEO during this remarkable era of Digital transformation, as ChainSys is on an exciting trajectory to support this shift and help customers design, implement, and deploy the latest innovations in enterprise data management. We are honored to have a great customer base, hard-working & multi-skilled employees across the globe, and a respectful position in all the markets where we compete.

Sundu Rathinam Founder/ CEO/ CTO, ChainSys



US Headquarters

USA - MICHIGAN
325 S. Clinton St., Suite 205
Grand Ledge, MI 48837

517-627-1173

US Solutions and Support HQ

USA - CALIFORNIA
CHAIN-SYS CORPORATION
500 Menlo Drive, Rocklin,
CA 95765

517-627-1173 Ext: 2

Europe Solutions & Support Center

EUROPE
CHAIN-SYS EUROPA B.V.
Jan Pieterszoon Coenstraat 7
The Hague 2595 WP
The Netherlands

Solutions & Development Center

MIDDLE EAST
CHAIN SYS MIDDLE EAST FZCO
2E 401, Fourth Floor,
DAFZA, Dubai, UAE
PO Box no. 371425

+971 042578847

Global Development Center

INDIA - MADURAI
CHAIN-SYS SOFTWARE EXPORTS
PRIVATE LIMITED
ELCOT IT Park, SEZ-2
Vadapalanji,
Madurai - 625 021

Global Center Of Excellence

INDIA - CHENNAI
CHAIN-SYS (INDIA) PRIVATE LIMITED
#85, Ponnamman Nagar,
Ayanambakkam,
Chennai - 600095

+91 (44) 69244100

www.chainsys.ai