



Transforming Oracle Data to Business Intelligence:

Creating Context-Rich Enterprise
Data with Smart Data Platform

Executive Summary

Oracle ERP and Fusion systems capture critical business transactions accurately and at scale. But transactional data is not automatically decision-ready. Technical structures show where data is stored, not what a metric means, who owns it, which rules apply, or whether it is fit for a specific decision.

When this business context remains scattered across documents, spreadsheets, and individual expertise, teams interpret the same data differently. Finance, sales, and operations produce conflicting metrics, analysts spend time reconciling results, and leaders question dashboards that cannot clearly show definitions, lineage, transformations, and data quality.

ChainSys Smart Data Platform closes both gaps at once. It attaches business definitions, lineage, and governance directly to your Oracle data, through a business glossary, automated lineage mapping, and policy enforcement that runs continuously instead of sitting in a document nobody opens. The result is data a person can trust and a system can interpret correctly, the same context doing both jobs.

This paper explains why Oracle data often stalls at “available” rather than becoming “understandable,” what context-rich enterprise data looks like in practice, and how it is built through cataloging, analytics, and security. It also explores how ChainSys supports a phased path from discovery to trusted business intelligence

The Trust Gap, By Numbers

29%

Only 29% of employees at companies with BI tools actually use them, per seven years of Gartner tracking. Structure without meaning doesn't get trusted, no matter how good the tool is.

42%

A 2026 survey of 505 data and analytics leaders found 42% believe their infrastructure is ready for AI and analytics, and name that same infrastructure as their biggest obstacle, in the same breath.

80%

80% of heavy data and analytics users, per a June 2026 Databox survey, have no single unified view of their data. Only 9% do, roughly the share of companies that have actually solved the problem this whitepaper is about.

Table of Contents

1. Introduction: The Oracle-to-BI Imperative	1
1.1 Why Oracle data is strategic but often underused	1
1.2 From dashboards to decisions: the shift to context-rich enterprise data	2
1.3 Introducing ChainSys Smart Data Platform	2
2. The Problem: Why Oracle Data Stalls	3
2.1 The hidden data model problem in Oracle ERP and Fusion	3
2.2 The context gap: technical metadata versus business meaning	3
2.3 Consequences for BI and AI initiatives	4
3. The Opportunity: Defining Context-Rich Enterprise Data	5
3.1 What is context-rich enterprise data?	5
3.2 Catalog versus a governed context layer	5
3.3 The business outcomes of getting context right	6
4. Reference Architecture: From Oracle Data to Trusted Insights	7
4.1 High-level architecture overview	7
4.2 Data flows and integration patterns	7
4.3 Integration with the existing analytics stack	8
4.4 Where ChainSys Smart Data Platform fits	9
5. Pillar I: Cataloging, Turning Oracle Metadata into Business Context	10
5.1 Automated discovery and profiling of Oracle assets	10
5.2 Business glossary and metadata enrichment	11
5.3 Lineage and continuous monitoring	11
5.4 Search and discovery	11
6. Pillar II: Analytics, Building Trusted, Decision-Grade Views	12
6.1 Consistent metrics and dimensions	12

6.2 Data quality, profiling, and observability	12
6.3 Visualization and reporting	13
6.4 Machine learning inside the analytics engine	13
6.5 Asking a question in plain language	14
7. Pillar III: Security, Governing Access, Use, and Compliance	15
7.1 Access control and identity	15
7.2 Identifying and validating sensitive data	16
7.3 Masking, encryption, and policy enforcement	16
7.4 Governed remediation	16
7.5 Audit-ready trails	17
8. Implementation Playbook: A Pragmatic Path to Value	18
8.1 Start with discovery and assessment	18
8.2 Prove value on real data before scaling	18
8.3 Roll out in phases	19
8.4 Manage change and train for adoption	19
8.5 Govern after launch	19
9. Measuring Impact: From Pilot to Program	19
9.1 What to track	19
9.2 Risk and compliance outcomes	20
9.3 Business value in practice	20
10. Future-Proofing: Why This Groundwork Pays Off Twice	21
10.1 The context gap affects BI and AI	21
10.2 Built to carry forward	21
10.3 A reusable enterprise asset	22
Conclusion: Make Context a Competitive Advantage	22
Appendix A: Capability Checklist for Smart Data Platform	23
Cataloging	23
Analytics	23
Security	23
Appendix B: Glossary	24

1. Introduction: The Oracle to BI Imperative

1.1 Why Oracle data is strategic but often underused


Oracle systems contain the operating record of an enterprise. A purchase order, supplier payment, inventory movement, employee change, or revenue transaction has to be handled correctly. That is what Oracle EBS and Oracle Cloud ERP are designed to do.

Analytical readiness is a different discipline. A data structure optimized to process a purchase order is not automatically suitable for answering a finance leader's question six months later about budget variance, supplier exposure, or working-capital risk. To answer that question reliably, teams need more than the stored values. They need a common definition, the relationships that connect the data, and evidence of how the metric was produced.

Without that context, BI initiatives tend to create more reports but not necessarily more confidence. Different teams build their own data extracts and KPI logic. The business then spends time reconciling dashboards instead of acting on them.

Oracle records what happened.

The business needs to know **what it means**



One metric. Three answers.

The quarter gets spent reconciling dashboards instead of acting on them

FINANCE

\$84.2M

SALES

\$84.2M

OPERATIONS

\$84.2M

1.2 From dashboards to decisions: the shift to context-rich enterprise data

A dashboard is a delivery mechanism. A decision requires confidence in the evidence behind it. That confidence comes from answering a few basic questions quickly:



What does this metric mean in our business?

Which Oracle and non-Oracle sources contributed to it?

What transformations, rules, and filters were applied?

Is the data current, complete, and approved for this use?

Who owns the definition and can approve a change?

This is the difference between technical metadata and business context. Technical metadata describes a table, column, schema, or data type. Business context explains how that information should be interpreted and used. Modern BI and AI-assisted analytics need both.

1.3 Introducing ChainSys Smart Data Platform

The ChainSys Smart Data Platform catalogs, analyzes, and secures enterprise data. It also handles data migration, integration, and master data management, all on one platform. It's been deployed across more than 500 implementations, including multiple Fortune 500 and Fortune 100 organizations.

For Oracle-centric enterprises, the platform provides the context layer between operational systems and decision-making tools. It is designed to complement existing Oracle and BI investments, not replace them. The platform helps organizations turn technical metadata into governed business knowledge that can be searched, validated, monitored, and applied across reporting and analytics.

2. The Problem: Why Oracle Data Stalls

2.1 The hidden data model problem in Oracle ERP and Fusion

Every Oracle implementation reflects the business it serves. Teams configure flexfields, add custom objects, encode approval logic, and establish local rules for processes that are unique to the organization. None of that configuration ships with a translation guide. It may work perfectly in the application while remaining difficult for an analyst or new data steward to interpret.

Even standard objects need business translation. A field such as STATUS_CODE may be documented technically, but its meaning can vary with a company's process, policy, workflow, and reporting rules. The value is available. Its business meaning is not always readily available.


2.2 The context gap: Technical metadata versus business meaning

Think of it as barcodes instead of labels. Oracle's schema gives every table and column a barcode: a technical identifier that IT can scan and look up. What it doesn't give anyone outside IT is the label, the plain-language explanation of what that field represents, which policy governs it, and who's accountable if it's wrong.

That distinction matters because the missing layer isn't stored anywhere that the machine can read. Process logic, business rules, security constraints, and the relationships between modules exist, but they exist as tribal knowledge and static documentation. As a result, every new BI initiative must rediscover information that the enterprise should already manage as a shared asset.

Why Data Stalls

Barcodes without Labels



2.3 Consequences for BI and AI initiatives

The first consequence is inconsistent KPIs. Finance, sales, and operations each build their own version of "revenue" or "active customer," and reconciling the differences becomes a recurring monthly tax on everyone's time. The issue is not necessarily that the source transactions are wrong. The issue is that the definitions and rules are not governed as a common enterprise standard.

The second is slower time-to-value on any BI or analytics project. Analysts must investigate data structures, validate joins, and ask subject-matter experts to confirm definitions before they can build a trustworthy view.

The third, and the one growing fastest, is when undefined data is used in advanced analytics or AI initiatives: without clear definitions and lineage. The outputs are only as trustworthy as the guesswork behind them, and low trust in analytics quickly becomes low trust in the outcome.

None of this is a data quality problem in the traditional sense. **The data is usually accurate. It's an understanding problem.** The understanding isn't fixed by cleaning data, it's fixed by attaching meaning to it.


Critical Business Consequences

When Business Context is Missing



3. The Opportunity: Defining Context-Rich Enterprise Data

Context-Rich Enterprise Data



3.1 What is context-rich enterprise data?


Context-rich enterprise data carries more than raw values. It includes four elements: a clear business definition, rules governing its use, relationships to other data, and a traceable history. Together, these elements help users trust a number without repeatedly validating it. They also give BI tools and governed systems the context needed to interpret data consistently.

3.2 Catalog versus a governed context layer

A data catalog answers two questions: what data exists and where. That is valuable, but not sufficient. It may identify a table such as AR_INVOICES, yet still fail to explain how the business defines “past due” under its credit policy or how that definition changes over time.

A governed context layer goes further. It connects definitions, lineage, ownership, policies, and access controls to the data and keeps that context current. The catalog shows that the data exists. The context layer explains what it means, who can use or change it, and whether it can be trusted.

Active Metadata Management, From Source System to Governed Context



3.3 The business outcomes of getting context right

Governed context supports faster, more confident decisions because users spend less time verifying numbers. Analytics and AI outputs become more reliable because definitions and lineage remain consistent across projects. Audit and compliance effort also falls because controlled access and a traceable history replace manual reconstruction after the fact.

Context Element	What It Answers	Why It Matters for BI
Business Definition	What does this metric, field, or dataset mean?	Prevents competing KPI interpretations.
Ownership and Stewardship	Who is accountable for accuracy and approval?	Creates a clear path for change and issue resolution.
Relationships	How do entities connect across modules and systems?	Reduces incorrect joins and incomplete analysis.
Lineage	Where did the data originate, and how was it transformed?	Makes insights explainable and audit-ready.
Quality Signals	Is the data complete, consistent, and current enough to use?	Helps users judge its fitness for purpose.
Security and Policy	Who can use the data, and under what conditions?	Protects sensitive information without blocking legitimate analysis.

4. Reference Architecture: From Oracle Data to Trusted Insights

4.1 High-level architecture overview

An effective architecture connects Oracle source systems to a governed data and metadata layer before information reaches analytics and BI tools. Sources may include Oracle E-Business Suite, Oracle Cloud ERP, PeopleSoft, JD Edwards, Siebel, and other enterprise applications within the broader data estate.

The architecture consists of four connected layers:



Source and connectivity layer:

Connects Oracle and non-Oracle systems through APIs, database connectors, file-based extracts, and event-driven integrations.



Metadata and context layer:

Captures technical metadata, business definitions, ownership, relationships, classifications, and end-to-end lineage.



Trust and governance layer:

Applies data-quality rules, access policies, stewardship workflows, monitoring, and audit controls.



Analytics and consumption layer:

Delivers trusted, context-rich data to existing BI platforms, dashboards, reporting tools, and downstream analytics applications.

Together, these layers create a traceable path from operational Oracle data to trusted business insight.


4.2 Data flows and integration patterns

Oracle data may move through batch extracts, APIs, database connections, events, or Change Data Capture, depending on the source system and analytical requirement. Change Data Capture helps keep metadata, lineage, and data profiles current as records and structures change.

As data moves through the architecture, it is:

- Discovered and cataloged
- Profiled for completeness, consistency, and accuracy
- Mapped to business definitions and ownership
- Connected through technical and business lineage
- Classified according to sensitivity and policy
- Monitored for changes and quality issues
- Prepared for governed analytical consumption

The architecture should accommodate structured records as well as unstructured content, including documents and attachments associated with Oracle business processes. This prevents relevant information from remaining disconnected from the wider governance and analytics environment.



4.3 Integration with the existing analytics stack

This architecture complements existing Oracle, data, and BI investments rather than replacing them. Governed data and metadata can be made available to current dashboards, reporting platforms, data warehouses, cloud data platforms, and analytical applications.

Security and policy controls operate across every layer. They determine who can discover, access, interpret, and use the data while preserving lineage and audit evidence. The result is an environment in which users can trace an insight back to its source, understand its business meaning, assess its quality, and use it with confidence.

4.4 Where ChainSys Smart Data Platform fits

The ChainSys Smart Data Platform brings this reference architecture together by connecting metadata discovery, business context, lineage, data profiling, governance, and analytics across Oracle and non-Oracle sources. It enriches technical metadata with definitions, ownership, relationships, quality signals, and policy controls.

The platform does not replace Oracle applications, data platforms, or BI tools. It provides the governed context layer between them, helping users find, understand, trace, secure, and use enterprise data with greater confidence. As systems and business rules change, this context can be monitored and kept current.


5. Pillar I: Cataloging, Turning Oracle Metadata into Business Context

5.1 Automated discovery and profiling of Oracle assets

Smart Data Platform connects to Oracle and other enterprise sources to extract technical metadata, including schemas, tables, columns, and relationships. Automated profiling then examines data structures, patterns, and content to assess quality and identify likely connections between entities. This gives teams a reliable starting point without depending on manual documentation.

In the illustration below, the extract adapter moves the data for profiling, tagging, lineage, and cataloging.

Data Catalog Component Architecture, from Source Endpoints to Catalog Handler



Business context also extends beyond structured tables. A complete view of supplier exposure, for example, may depend on both purchase orders and contract documents. Smart Data Platform profiles text, tabular files, forms, and images through the same connector framework. A document module extracts and classifies this content, then adds the results to the catalog. This closes a common governance gap by bringing relevant documents into the same context layer as transactional data.

5.2 Business glossary and metadata enrichment

The business glossary turns technical fields into understandable business assets. It adds plain-language definitions, ownership, domain, tags, attributes, and relationships to registered tables and fields. A column such as **CUSTOMER_ID** no longer stands alone; users can see what it means, how it relates to other data, and who is accountable for it.

The Platform also provides workflows to register, review, and publish glossary objects across the organization. Once a steward approves a definition, it becomes a reusable enterprise standard rather than another isolated document.

AI supports this work by helping teams:

- Profile unfamiliar or custom applications
- Predict likely attribute values during registration
- Suggest relationships and lineage connections

These suggestions remain subject to steward review. Approval workbenches, tag approvals, and access-request workflows keep accountability with the business while reducing manual effort.

5.3 Lineage and continuous monitoring

Automated lineage and entity-relationship mapping show how data moves and changes from source to report. Users can trace a dashboard metric through its transformations and back to the originating system, making the result easier to validate and explain.

Continuous monitoring keeps this map current as schemas, pipelines, and configurations change. It also helps teams identify downstream impact before a change reaches business reporting.

5.4 Search and discovery

Indexed search allows business and technical users to find data by business term, tag, or entity without knowing the underlying table name. Data citizens can also view, comment on, and approve suggested metadata, keeping the catalog active and current.

A built-in conversational assistant extends discovery through plain-language questions. It searches governed metadata, including table and column names, descriptions, tags, and relationships, then maps business terms to the relevant schema elements. Users can find the right data even when they do not know how the source system is structured.


This capability depends on the context already established in the catalog. Definitions, tags, relationships, and access controls give the assistant a governed foundation for interpretation. Without them, it has little reliable context on which to reason.

6. Pillar II: Analytics, Building Trusted, Decision-Grade Views

6.1 Consistent metrics and dimensions

Consistent metrics start with a shared source. Datasets, dimensions, and KPIs are built on the governed catalog, so the definitions that make data understandable also shape what appears in reports. Data modelers and calculated tables allow teams to define a metric once and reuse it, instead of rebuilding the logic in separate extracts.

The result is a shared analytical foundation across finance, sales, operations, and other functions. Local analysis remains possible, but any departure from the enterprise definition becomes visible and deliberate.



6.2 Data quality, profiling, and observability

Data-quality rules, scores, and measurements are linked directly to metadata. Users can assess accuracy, completeness, consistency, and timeliness in the context of the data they plan to use, rather than consulting a separate quality report after the fact.

Quality thresholds can reflect the purpose of each dataset. An executive KPI may require stricter controls than exploratory analysis, but both remain transparent and measurable.

6.3 Visualization and reporting


Teams can present governed data in more than 25 visualization types, including:

- Hierarchies, pivot tables, and heatmaps
- Geographic maps and Gantt charts
- Bar, treemap, funnel, bubble, and ring charts

Teams can create dashboards from templates, group them by purpose, and assign them by role. Reports and dashboards can also be scheduled for distribution and exported for use in other tools. The value lies not only in presentation, but in ensuring that every view is built on governed definitions and access controls.

Sample Chart Output

Growth Report




6.4 Machine learning inside the analytics engine

The learning engine supports supervised, unsupervised, and reinforcement learning, together with natural language processing. These capabilities extend analysis beyond reporting what happened to identifying patterns, estimating likely outcomes, and informing possible actions.

Most importantly, the models operate on governed, context-rich data. Definitions, lineage, and quality signals reduce the ambiguity that weakens analytical and AI outcomes.

Analytics and Visualization Engine Architecture, Including the Learning Engine



6.5 Asking a question in plain language

The conversational assistant can answer questions at three controlled levels:



Metadata

Searches names, descriptions, tags, and relationships to identify relevant data assets.



Data lake

Queries a configured source only when Enterprise Data Management is enabled and the lake is synchronized. Otherwise, it uses datasets already available in the lake.




Dataset

Answers only from the tables within a dataset for which an administrator has enabled the AI feature.

Administrators define the scope, not the person asking the question. Existing row- and column-level permissions still apply, and masked fields remain masked. Access becomes easier without expanding a user's authorized reach.

As ChainSys documentation explains, the dataset-level approach "ensures controlled data access, provides dataset-specific insights, improves governance and security."

Plain language. Controlled reach.



7. Pillar III: Security, Governing Access, Use, and Compliance

7.1 Access control and identity

Role-based access control defines who can view, edit, or approve data down to the row and column level. Support for SAML, OAuth 2.0, and JWT allows the platform to integrate with identity providers such as Active Directory and Okta. Access therefore follows the organization's existing identity model and least-privilege policies.

7.2 Identifying and validating sensitive data

Structured and unstructured sources are scanned for personally identifiable information (PII). Organizations can configure patterns, rules, lookup values, PII categories, and exclusions to reflect their own risk definitions.

Exclusions are important in large Oracle estates, where broad scans can create false positives and overwhelm review teams. Configurable criteria keep findings relevant. Business users then validate proposed classifications, with bulk approval and rejection available for scale. The system identifies likely risks; accountable people confirm it.

7.3 Masking, encryption, and policy enforcement


Confirmed sensitive data can be masked at the column level, while encryption protects data at rest and in transit. Automated policy enforcement applies governance rules as data is accessed and used. Continuous monitoring supports regulatory requirements, including GDPR and CCPA, by surfacing exposure when it occurs rather than during a periodic review.

7.4 Governed remediation

Detection alone does not reduce risk. Remediation profiles, configuration-driven PII remediation, approval controls, and execution tracking turn findings into action. Teams can see what was planned, approved, completed, or unsuccessful.

Because remediation changes real data, approval separates the plan from execution and every run retains a status trail. The platform automates discovery and execution while keeping material decisions with authorized people.

Finding only Sensitive Data Is the First Step; Governed Remediation Completes the Process



7.5 Audit-ready trails

Every change across classification, approval, access, and remediation is recorded. On-demand reporting gives governance, risk, and compliance teams direct evidence of what happened, who approved it, and what action followed. This replaces time-consuming reconstruction with a continuous, reviewable record.

Security and Compliance Capability Map

Capability	What It Covers
Role-based access control	Row- and column-level access based on least-privilege principles
Single sign-on	SAML, OAuth 2.0, and JWT integration with Active Directory and Okta

Capability	What It Covers
Automated PII detection	Detection across structured and unstructured data
Data masking and encryption	Protection for data at rest and in transit
Automated policy enforcement	Governance rules applied during data use
Continuous audit trail	On-demand reporting for governance, risk, and compliance requirements
Platform certifications	SAP Certified, Oracle Certified, Salesforce Approved, IBM Ready, and Microsoft Ready

8. Implementation Playbook: A Pragmatic Path to Value

8.1 Start with discovery and assessment

Begin with visibility, not an enterprise-wide rollout. Connect priority systems, discover metadata and relationships, profile quality, and identify sensitive data. Combine these findings with stakeholder interviews to align the initiative with business priorities and regulatory needs. The outcome should be a clear estate map and a focused implementation plan.

8.2 Prove value on real data before scaling

Run a focused proof of value against a representative subset of Oracle data. Demonstrate outcomes stakeholders can verify directly, such as:



Faster data discovery



Automated stewardship workflows



Traceable end-to-end lineage

The objective is evidence from real data, not a conceptual demonstration.

8.3 Roll out in phases

Scale in controlled waves, reusing approved configurations and templates:



Onboard
source systems



Extract
Metadata



Profile and
enrich data



Onboard
users

This approach limits risk, accelerates early value, and creates a repeatable pattern for each new domain.

8.4 Manage change and train for adoption

Role-based training should address the different needs of business users, data stewards, analysts, and administrators. Practical workshops and ongoing support help teams incorporate governed data into everyday work. Adoption, not installation, determines whether the platform delivers value.

8.5 Govern after launch

Stewardship roles, curation workflows, usage monitoring, and change controls must continue after deployment. This keeps definitions, lineage, and ownership aligned as systems and business rules evolve, preventing the catalog from becoming another static repository.

9. Measuring Impact: From Pilot to Program

9.1 What to track

Measure both coverage and adoption:

Coverage

The proportion of priority Oracle assets with approved definitions, named owners, quality scores, sensitive-data classifications, and mapped lineage.

Adoption

Successful catalog searches, use of certified datasets, reuse of governed KPIs, and reduced duplication of analytical work.

These indicators show whether the organization is building context and whether people are using it.

9.2 Risk and compliance outcomes

Track improvements in policy compliance, approval turnaround, sensitive-data remediation, and audit-response time. Continuous audit trails and role-based access should make evidence available on demand, reducing the effort required to support GRC reviews.

9.3 Business value in practice

An Ohio-based global provider of power, cooling, and IT infrastructure solutions illustrates the potential. Following a merger and a move from legacy mainframes to a hybrid Oracle EBS and Oracle Cloud environment across APAC, the Americas, and EMEA, the company struggled to locate and trust data across more than 30 applications.

Within 12 weeks, Smart Data Platform centralized metadata from those applications in a single catalog, established an enterprise business glossary, and introduced automated PII detection. Business users gained a searchable view of data that had previously required duplicated manual effort to assemble and validate.

Across ChainSys platform implementations, organizations typically report 25 to 30 percent time and cost savings in the first implementation, increasing to as much as 50 percent as adoption matures. These gains come from reusable templates and repeatable processes rather than rebuilding each initiative from the ground up.

10. Future-Proofing: Why This Groundwork Pays Off Twice

10.1 The context gap affects BI and AI

The same missing context that makes users question a dashboard also limits AI. A model can access a table, but it cannot safely infer the organization's definition, policy, ownership, or acceptable use. Structure without meaning produces fragile results, regardless of the technology consuming it.

Stage	What It Looks Like
Ad hoc	Definitions and lineage remain in people's knowledge and legacy documents.
Cataloged	Metadata is centralized and searchable but is not yet tied to enforced business rules.
Governed	Glossary terms, lineage, ownership, and access policies are connected to data and maintained continuously.
Context-rich	Business meaning, lineage, quality, and governance are complete enough to support trusted decisions without repeated reconciliation.
AI-ready	The same governed context supports predictive, prescriptive, and AI-driven use cases without a separate cleanup program.

10.2 Built to carry forward

The context built for BI does not need to be rebuilt when the organization expands its use of AI-driven analytics. The same glossary, lineage, quality signals, and governance controls provide the foundation required for predictive and prescriptive analysis. The organization extends an existing capability instead of starting a second program.

10.3 A reusable enterprise asset

Treat context as a durable enterprise asset, not a project deliverable. As the Oracle estate changes and new human or automated consumers use its data, the context layer allows each use case to begin with shared definitions, visible lineage, and governed access rather than another round of discovery and reconciliation.

Conclusion: **Make Context a Competitive Advantage**

Oracle systems hold the operational record of the enterprise. Turning that record into trusted business intelligence requires a governed layer of definitions, lineage, quality, ownership, and access controls.

The ChainSys Smart Data Platform brings those capabilities together across cataloging, analytics, and security. Organizations can start with a focused business domain, prove value on real data, and scale through reusable patterns. With more than 500 implementations, Fortune 100 organizations, ChainSys offers the experience to support an Oracle-centered rollout.

The gap between available data and confident decisions is solvable. Start with one high-value domain, such as order-to-cash or supplier spend, and see what it takes to make that data trusted.

To see how the ChainSys Smart Data Platform does this on Oracle data, **request a demo** at **chainsys.ai/demo-request**.

Appendix A: Capability Checklist for Smart Data Platform



Cataloging

- Automated metadata harvesting
- AI-assisted profiling
- Business glossary
- Metadata tagging
- Attribute and relationship mapping
- Automated lineage and entity relationship mapping
- Advanced indexed search
- Continuous monitoring



Analytics

- Consistent metrics and dimensions built on governed metadata
- Integrated data quality scoring
- 25-plus chart types and customizable dashboards
- Scheduled reporting
- AI-assisted descriptive, predictive, and prescriptive analysis



Security

- Role-based access control down to row and column level
- Single sign-on (SAML, OAuth 2.0, JWT)
- Automated PII detection
- Data masking and encryption at rest and in motion
- Automated policy enforcement
- Continuous audit trail and GRC reporting
- SAP, Oracle, Salesforce, IBM, and Microsoft certifications

Appendix B: Glossary

Active metadata management

ChainSys's term for metadata that gets applied rather than only recorded. A definition or a policy held this way is enforced when the data is used, instead of sitting in a document nobody opens.

Business glossary

A plain-language record of what each field means to the business, who owns it, and which domain it belongs to, attached to the actual tables and columns.

CCPA

The California Consumer Privacy Act. State law giving California residents rights over how their personal data is collected and used.

Change Data Capture

A method of spotting changes in a source system as they happen, so a catalog stays current without waiting for a full manual refresh.

Context-rich data

Data that carries four things alongside its raw values: a readable business definition, the rules governing its use, its relationships to other data, and a visible trail of where it came from.

Data catalog

An inventory of what data exists and where it sits. It answers the location question. On its own it does not carry business meaning or enforce anything.

Data citizen

A business user who works with cataloged data without being a technical specialist. In Smart Data Platform they can view, comment on, and approve suggested metadata.

Data lake

A store holding data brought together from several sources, in varied structures, without forcing everything into one schema first.

Data lineage

The record of where a piece of data came from, how it moved, and what changed it on the way. It is what lets someone trace a number on a dashboard back to its source.

Data masking

Hiding or replacing the real values in sensitive fields, so the data stays usable without exposing what it contains.

Data profiling

Examining data to learn its structure, patterns, and content. This is how quality problems and relationships between tables get spotted.

Data steward

The person accountable for a set of data. In the Smart Data Platform, the steward approves or rejects the definitions, tags, and relationships the system suggests.

Dataset

A defined collection of tables and fields prepared for analysis. An administrator can switch the AI assistant on for one dataset at a time, which limits what it is able to reach.

Descriptive, predictive, and prescriptive analysis

Three levels of analysis. Descriptive reports what happened. Predictive estimates what is likely to happen next. Prescriptive suggests what to do about it.

Enterprise Data Management (EDM)

The practice of defining, integrating, governing, and maintaining an organization's data across all its systems. In this paper, it also refers to a setting at the data lake level. Switched on, the chat assistant can reach a configured source connection. Switched off, it answers only from datasets already held in the lake.

Flexfield

An Oracle mechanism for adding custom fields to standard objects without altering the core schema. Useful for tailoring the system, and a common source of undocumented data.

GDPR

The General Data Protection Regulation. European law governing how personal data is collected, stored, and used.

Governed context layer

The layer this paper argues for. It attaches definitions, lineage, and enforced governance to data continuously, so the meaning stays current and access stays controlled.

GRC

Governance, risk, and compliance. The function responsible for proving an organization follows its own rules and the regulations that apply to it.

Identity provider

A system that holds user identities and confirms who someone is, such as Active Directory or Okta. Other applications trust it rather than keeping separate logins of their own.

Machine learning

Models that improve from data rather than from fixed rules. Supervised learning trains on labeled examples. Unsupervised learning finds patterns without labels. Reinforcement learning improves by trying actions and learning from the outcome.

Master data management (MDM)

Maintaining one agreed version of core business records such as customer, supplier, and product.

Metadata

Information about data rather than the data itself. Technical metadata covers table names, column types, and schemas. Business metadata covers what a field means, who owns it, and which rules apply to it. Oracle supplies the first. The second has to be built.

Metadata tagging

Labeling data by business concept, so people can search by what something means rather than by table name.

Natural language processing (NLP)

Technology that lets a system work with ordinary human phrasing. It is what allows a question typed the way a colleague would ask it to be understood.

Oracle Cloud ERP

Oracle's cloud enterprise resource planning suite, also referred to as Oracle Fusion.

Oracle E-Business Suite (EBS)

Oracle's long-established on-premise enterprise resource planning suite.

Personally identifiable information (PII)

Data that identifies a person on its own, or when combined with something else. Names, email addresses, phone numbers, and government identifiers all count.

Remediation

The step after detection. Fixing or protecting the sensitive data a scan has found, rather than only producing a list of it.

Role-based access control (RBAC)

Granting access by a person's role rather than person by person. In the Smart Data Platform, this reaches down to specific rows and columns.

Single sign-on (SSO)

One set of credentials giving access to several applications. The Platform supports it through the SAML, OAuth 2.0, and JWT standards.

Smart Data Platform

ChainSys's platform for managing enterprise data from end to end. It covers data migration, integration, archival, and reconciliation; data quality management and master data governance; and data cataloging, analytics, visualization, masking, and compliance.

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

Authors



Amarpal Nanda

President of EDM
amarpal.nanda@chainsys.ai



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