

ORACLE WHITEPAPER

ORACLE DATA INTEGRATOR ENTERPRISE EDITION

ORACLE DATA INTEGRATOR ENTERPRISE EDITION



Oracle Data Integrator Enterprise Edition delivers high-performance data movement and transformation among enterprise platforms with its open and integrated E-LT architecture and extended support for Big Data. Oracle Data Integrator Enterprise Edition is critical to leveraging data integration initiatives on-premise or in the cloud, such as Big Data management, Service Oriented Architecture and Business Intelligence. An easy-to-use user interface combined with a rich extensibility framework helps Oracle Data Integrator Enterprise Edition improve productivity, reduce development costs and lower total cost of ownership among data-centric architectures. Oracle Data Integrator Enterprise Edition is fully integrated with Oracle Fusion Middleware, Oracle GoldenGate, Oracle Database, Oracle Big Data Appliance and Exadata to put data at the center of your enterprise. Oracle Data Integrator Enterprise Edition is open and standards-based to work with 3rd party applications as well as Oracle's applications.

KEY FEATURES

Out-of-box integration with databases, Hadoop, ERPs, CRMs, B2B systems, flat files, XML, JSON, LDAP, JDBC, ODBC

Knowledge module framework for extensibility

Powerful data transformation for heterogeneous database and Big Data infrastructures

Rich ETL for Oracle databases including Oracle Exadata, with complex dimension and cube loading support

Integrates with Oracle GoldenGate for real-time data warehousing

Metadata-driven data lineage and impact analysis with Oracle Enterprise Metadata Management

Integrates with Oracle Enterprise Data Quality for advanced profiling,

About Oracle's Data Integration Solutions

Oracle is a leader in the Data Integration market, with the industry's most comprehensive fully-integrated offering for data integration, including Oracle Data Integrator Enterprise Edition, Oracle GoldenGate, Oracle Enterprise Data Quality, Oracle Big Data Preparation Cloud Service and Oracle Enterprise Metadata Management. Oracle's Data Integration solutions provide continuous access to timely, trusted, and heterogeneous data across the enterprise to support both analytical and operational data integration.

What's new in 12c?

New capabilities included in Oracle Data Integrator Enterprise Edition 12c include:

- Superior productivity with a new flow-based declarative user interface. The enhanced user experience is simple yet powerful and comprehensive; the ability to seamlessly reuse mapping logic speeds up development and overall time to solution.
- Improved parallelism functionality increases Oracle Data Integrator's already high performance architecture.
- High performance E-LT capabilities integrated with Oracle GoldenGate to enable faster and more efficient loading and transformation of real-time data into a data warehouse or Big Data system. Customers can now easily configure and deploy real-time data warehousing solutions without impacting source systems or batch window dependencies.

cleansing, matching and data governance needs.

KEY BENEFITS

Native big data Support

Leading Performance

Improved Productivity

Flexibility

Low Total Cost of Ownership

Expanded Heterogeneity to big data sources and targets

UNIQUE BENEFITS

Cuts hardware costs through improved utilization and high-performance data integration

Unifies data-based, event-based, and service-based integration in a single solution

Lowers total cost of ownership and by maximizing designer productivity

Open, 100% Java, SOA, Standards based

Supports heterogeneous sources, targets and applications

RELATED PRODUCTS AND SERVICES

Oracle Data Integrator Enterprise Edition delivers unique next-generation, extract load and transform (E-LT) technology that improves performance, reduces data integration costs and improves productivity.

RELATED PRODUCTS

Oracle Data Integrator Enterprise Edition Big Data Options

Oracle GoldenGate

Oracle Enterprise Data Quality

Oracle Enterprise Metadata Management

Oracle Exadata

Oracle Big Data Appliance

Oracle Business Intelligence

- Migration utility and improved interoperability with Oracle Warehouse Builder, providing the tools and foundation for customers to migrate to Oracle's strategic Oracle Data Integrator platform.
- Direct integration into Oracle Enterprise Manager Cloud Control 12c for improved IT productivity in monitoring Oracle technology and applications from one central location.
- Extended big data support in unifying integration tooling via industry standards across unstructured/semi-structured and structured data.
- Release management capabilities to allow customers to better manage their code as well as their development, testing and production environments. This includes a deep integration with Apache Subversion.

Oracle Data Integrator for Big Data

Oracle Data Integrator brings critical big data integration capabilities to enterprises. It brings heterogeneity to Big Data by extending support to Hadoop standards. Through its decoupled design and implementation principles, ODI mappings seamlessly can switch between underlying Big Data technologies. Currently ODI for Big Data natively supports Spark, Spark Streaming, Hive, Kafka, Cassandra, HBase, Sqoop and Pig. ODI also supports Apache Oozie scheduling, in addition to the traditional ODI agent to schedule and orchestrate big data mappings, procedures, packages and scenarios. Out of the box Knowledge Modules provide quick implementation and time to value to big data projects. [Click here to learn more about Oracle Data Integrator for Big Data.](#)

Simplifying the Complexity of Data Integration

Oracle Data Integrator Enterprise Edition addresses multiple enterprise data integration needs.

- **Data Warehousing and Business Intelligence**—by executing high-volume, high-performance loading of data warehouses, data marts, On Line Analytical Processing (OLAP) cubes, and analytical applications, Oracle Data Integrator transparently handles incremental loads and slowly changing dimensions, manages data integrity and consistency, and analyzes data lineage.
- **Big Data**—by providing prebuilt integration with big data technologies, such as Spark, Pig, HDFS, Hive, HBase or Sqoop, businesses are able to leverage additional sources of data previously too large and unwieldy to gain benefits from. Oracle delivers integration for big data technologies leveraging a metadata driven process while enabling data integration resources to easily manage how big data is extracted, loaded, and transformed.
- **Service-Oriented Architecture (SOA)**—by calling on external services for data integration and by deploying data services and transformation services that can be seamlessly integrated within an SOA infrastructure. Oracle Data Integrator's architecture additionally provides support for high-volume, high-performance bulk data processing to an existing service-oriented architecture.
- **Master Data Management (MDM)**—by providing a comprehensive data synchronization infrastructure for customers who build their own data hubs, work with packaged MDM solutions, or coordinate hybrid MDM systems with integrated SOA process analytics and Business Process Execution Language (BPEL) compositions.
- **Migration**—by providing efficient bulk load of historical data (including complex transformations) from existing systems to new ones. Oracle GoldenGate then seamlessly synchronizes data for as long as the two systems coexist, and Oracle Data Integrator continues to complement as needed for transformations.

Oracle SOA Suite
 Oracle Database
 Oracle Data Warehousing
 Oracle Master Data Management

E-LT Architecture for High Performance

Oracle Data Integrator Enterprise Edition's Extract, Load, Transform (E-LT) architecture leverages disparate relational database management systems (RDBMS) or Big Data engines to process and transform the data. This approach optimizes performance and scalability and lowers overall solution costs.

Instead of relying on a separate, conventional ETL transformation server, Oracle Data Integrator Enterprise Edition's E-LT architecture generates native code for disparate RDBMS or big data engines (SQL, HiveQL, PySpark, Pig Latin or bulk loader scripts, for example). The E-LT architecture extracts data from the disparate sources, loads it into a target, and executes transformations using the power of the database or Hadoop. By leveraging existing databases and big data infrastructures, Oracle Data Integrator Enterprise Edition provides unparalleled efficiency and lower cost of ownership. By reducing network traffic and transforming data in the server containing the target data, the E-LT architecture delivers the highest possible performance.

Heterogeneous Support

Oracle Data Integrator Enterprise Edition provides heterogeneous support for 3rd party platforms, data-sources, and data warehousing appliances. While Oracle Data Integrator Enterprise Edition leverages optimizations for Oracle Database to perform E-LT data movement, transformation, data quality and standardization operations, Oracle Data Integrator Enterprise Edition is fully optimized for mixed technologies including: sources, targets and applications, etc.

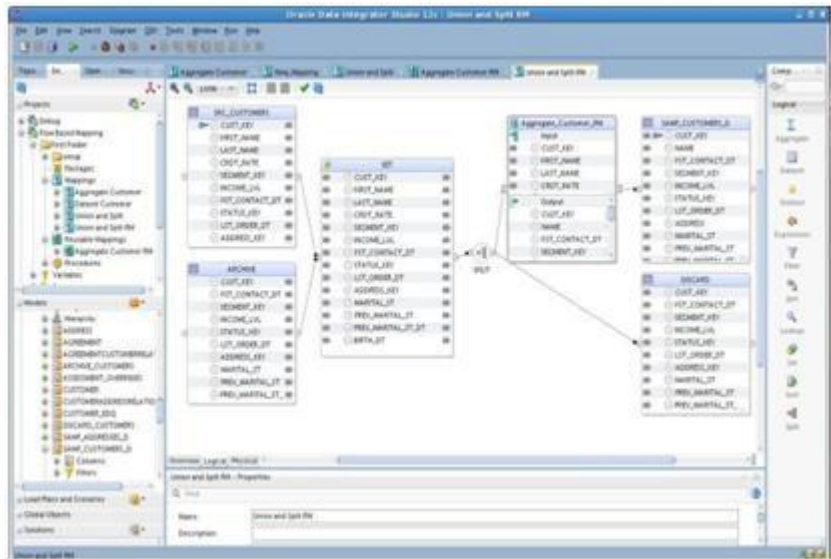


Figure 1 - Oracle Data Integrator Enterprise Edition Studio

High-Productivity Designer Paradigm for Data Integration

Oracle Data Integrator Enterprise Edition 12c introduces a new declarative flow-based user interface for enhanced user experience and productivity. The new user interface

combines the simplicity and ease-of-use of the declarative approach with the flexibility and extensibility of configurable flows. This blend simplifies common data integration design and deployment use cases, shortening implementation times. Data integration designers describe source and target data formats and data integration processes. The business user or the developer can focus on describing what to do, not how to do it. Oracle Data Integrator Enterprise Edition generates, deploys and manages the code required to implement those processes across the various source and target systems.

The paradigm is also enriched with the ability to seamlessly reuse mapping logic during development, giving developers a simpler and more efficient technique for providing solutions to their completion.

Knowledge Modules Provide Flexibility and Extensibility

Knowledge Modules are at the core of the Oracle Data Integrator Enterprise Edition's architecture. They make all Oracle Data Integrator Enterprise Edition processes modular, flexible, and extensible.

Knowledge Modules implement the actual data flows and define the templates for generating code across the multiple systems involved in each data integration process. Knowledge Modules are generic, because they allow data flows to be generated regardless of the transformation rules. At the same time, they are highly specific, because the code they generate and the integration strategy they implement are explicitly tuned for a given technology. Oracle Data Integrator Enterprise Edition provides a comprehensive library of Knowledge Modules, which can be tailored to implement existing best practices ranging from leveraging heterogeneous source and/or target systems, to methodologies for highest performance, for adhering to corporate standards, or for specific vertical know-how. By helping companies capture and reuse technical expertise and best practices, Oracle Data Integrator Enterprise Edition's Knowledge Module framework reduces the cost of ownership. It also enables metadata-driven extensibility of product functionality to meet the most demanding data integration challenges.

Enterprise Deployment for High Availability and Scalability

Oracle Data Integrator Enterprise Edition integrates to Oracle Fusion Middleware as a platform. Oracle Data Integrator Enterprise Edition provides its run-time components as Java EE applications, enhanced to fully leverage the capabilities of the Oracle WebLogic and Oracle Coherence. Oracle Data Integrator Enterprise Edition components include exclusive features for Enterprise-Scale Deployments, high availability, scalability, and hardened security.

High-Availability (HA) and Scalability is fully supported via clustered deployments for Java EE components. Oracle Data Integrator Enterprise Edition components deployed in WebLogic Server benefits from the capabilities of clustering for scalability, including JDBC connection pooling and load balancing. In addition to the cluster-inherited HA capabilities, the run-time agent also supports a connection retry mechanism to transparently recover sessions running in repositories that are stored in HA-capable database engines such as Oracle RAC.



Figure 2: Enterprise Manager Cloud Control Integration

Unified Administration and Management

Oracle Data Integrator Enterprise Edition simplifies complex data-centric deployments by improving visibility and control for with a unified set of management interfaces.

The Oracle Data Integrator Enterprise Edition Console leverages the Oracle Application Development Framework (ADF) and Ajax Framework for a rich user experience. Using this console, production users can set up an environment, export and import the repositories, manage run-time operations, monitor the sessions, diagnose the errors, browse design-time artifacts and generate lineage reports.

In addition, this interface integrates seamlessly with the Oracle Enterprise Manager Fusion Middleware Control Console and allows administrators to monitor from a single screen not only their data integration components but their other Fusion Middleware components as well.

To help maximize the value of Oracle Data Integrator, Oracle offers the Management Pack for Oracle Data Integrator. This plugin leverages Oracle Enterprise Manager Cloud Control's advanced management capabilities, to provide an integrated and top-down solution for your Oracle Data Integrator environments by providing a consolidated view of your entire Oracle Data Integrator infrastructure enabling users to monitor and manage all of their components centrally from Oracle Enterprise Manager Cloud Control.



CONTACT US

For more information about Oracle Data Integrator Enterprise Edition, visit oracle.com/goto/dataintegration or call +1.800.ORACLE1 to speak to an Oracle representative.

CONNECT WITH US

 blogs.oracle.com/oracle

 facebook.com/oracle

 twitter.com/oracle

 oracle.com

Copyright © 2016, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. 1015



ÜBER HUNKLER

Das Karlsruher Systemhaus HUNKLER wurde 1988 erster offizieller Partner von Oracle in Deutschland. Ein Team von rund 20 Mitarbeitern unterstützt Kunden aus Industrie, öffentlicher Verwaltung, Gesundheits- und Finanzwesen mit Beratung, Lösungsentwicklung und Managed Services.

Im Fokus von HUNKLER stehen leistungsfähige, wirtschaftliche Infrastrukturen für Oracle-datenbanken mit den Schwerpunkten Hochverfügbarkeit, Ausfallsicherheit und Zero Downtime Migration. Die integrierten Komplettlösungen der Produktfamilie Oracle Engineered Systems sowie der Datenbank-/Anwendungsbetrieb in der Oracle Cloud sind weitere Themenfelder, die das Unternehmen umfassend abdeckt.

HUNKLER
GmbH & Co. KG

Hauptsitz Karlsruhe

Bannwaldallee 32, 76185 Karlsruhe
Tel. 0721-490 16-0, Fax 0721-490 16-29

Geschäftsstelle Bodensee

Fritz-Reichle-Ring 6a
78315 Radolfzell
Tel. 07732-939 14-00, Fax 07732-939 14-04
info@hunkler.de, www.hunkler.de