

Oracle®

OFS MMG Application Pack Maintenance Release



Release 26.1.0

G52232-03

April 2026

ORACLE®

Oracle OFS MMG Application Pack Maintenance Release, Release 26.1.0

G52232-03

Copyright © 2014, 2026, Oracle and/or its affiliates.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software, software documentation, data (as defined in the Federal Acquisition Regulation), or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software," "commercial computer software documentation," or "limited rights data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

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

Intel and Intel Inside 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, Epyc, and the AMD logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Contents

- 1 Description
- 2 How to Apply this Release?
- 3 New Features, Enhancements and Limitations in this Release

1

Description

OFS MMG APPLICATION PACK RELEASE V26.1.0 (ID 39206590).

This release of the OFS MMG Application Pack is cumulative of all enhancements and bug fixes completed since the v26.1.0 release.

For more information on MMG related documents, see the following link: https://docs.oracle.com/cd/F40412_01/get_started.htm

2

How to Apply this Release?

For instructions on how to install this OFS MMG Application Pack Release, see the [Oracle Financial Services Model Management and Governance Application Pack Installation and Configuration Guide](#), Release 26.1.0.

3

New Features, Enhancements and Limitations in this Release

The OFS MMG 26.1.0 release introduces strategic architectural updates that are designed to enhance system governance, operational resilience, and analytical precision.

New Features

Monitoring and Observability

Consolidated Health and Performance Monitoring

- **Enhancement:** The legacy **Environment Health** tab has been replaced with a Prometheus-backed **Consolidated Health Service**. This introduces Current and History modes with detailed views into JVM metrics, JDBC connections, and Database Tablespace, delivering continuous, real-time telemetry and addressing previous data accuracy gaps.
- **Business Impact:** Establishes a single source of truth for environment health, enabling faster root-cause analysis, proactive resource management, and accurate historical trend analysis to help prevent application downtime.

Real-Time Pipeline Resource Monitoring

- **Enhancement:** The Model Pipeline Execution Summary now supports optional, high-granularity live tracking of CPU utilization, memory consumption, active threads, and Process IDs (PIDs) per session. Users can manually refresh telemetry or close sessions directly from the summary view.
- **Business Impact:** Provides immediate operational oversight to detect and terminate resource-heavy processes, prevent pipeline bottlenecks, and optimize compute costs.

Data and Pipeline Management

Dataset Transformation and DateTime Engineering

- **Enhancement:** The dataset management interface now supports full transformation lifecycle management — add, edit, delete, revalidate, and reorder — without switching context. Includes:
 - In-view application and preview of transformation logic against live dataset records.
 - Automated DateTime normalization across diverse formats for pipeline consistency.
 - Dedicated Encode Datetime functionality for ML-ready temporal feature engineering.
- **Business Impact:** Accelerates data preparation by enabling iterative transformations with direct data visibility, eliminating manual overhead and reducing error rates in time-series data cleaning.

Cross-Tenancy Data Porting and Enhanced Discovery

- **Enhancement:** Supports seamless data movement between Production and Sandbox tenancies. Key improvements:
 - Expanded schema and table discovery — no longer limited to mapped schemas only.
 - Selective data loading for granular control over what data is ported.

- Wider 75% viewport drawer with improved filtering, sorting, and clearer status messaging.
- **Business Impact:** Reduces time and effort to prepare test environments, improves accuracy navigating complex schema lists, and accelerates development cycles.

Service Interruption Handling

- **Enhancement:** The model pipeline now auto-tracks service disruptions. If a restart occurs during an active run, execution status automatically transitions to Interrupted — eliminating stale In-Progress states.
- **Business Impact:** Improves operational transparency and data integrity, enabling teams to quickly identify and re-trigger interrupted processes to maintain project timelines.

Scheduling and Automation

Dynamic Scheduler Input Configuration

- **Enhancement:** The Scheduler UI now supports dynamic, model-aware execution inputs. Parameter Sets display by default, available parameters auto-refresh on model change, and a standardized selection priority (**Parameter Sets** -> **Optional Params** -> **Execution Params**) ensures consistent, backward-compatible behavior.
- **Business Impact:** Eliminates manual parameter entry, reduces configuration errors, and ensures consistent model behavior across environments with automated priority logic.

Infrastructure and Integrations

Multi-Platform Git Integration and Validation

- **Enhancement:** The Git integration framework now supports GitHub, GitLab, and Bitbucket — removing the previous restriction to internal repositories. Includes automated repository URL validation and mandatory connection testing before any remote operations.
- **Business Impact:** Enables teams to use preferred version control providers; proactive URL validation minimizes failed deployments and broken syncs, saving developer time.

Kafka Topic and Cluster Management

- **Enhancement:** The Kafka management interface adopts a drawer-based interaction model for creating, editing, and viewing topics. Key updates:
 - Conda environment values can now be stored directly within topic records.
 - Support for comma-separated host:port bootstrap server lists, enabling multiple cluster nodes per topic.
- **Business Impact:** Drawer-based flows improve administrative productivity; multi-node bootstrap server support introduces high availability and fault tolerance, eliminating single points of failure in real-time data streaming.

Workspace-Aware Notification Management

- **Enhancement:** The Data Population notification framework is now a globally managed, workspace-aware system. It automatically seeds the default mapped user group per workspace, supports dynamic recipient resolution via a searchable list, and enforces workspace-scoped governance on all alerts.
- **Business Impact:** A set-and-forget approach that ensures critical alerts reach the correct stakeholders automatically — eliminating manual routing and improving organizational governance.

ML Platform and Model Management

AutoML Training Observability and Catalog Integration

- The Model Catalog now features a dedicated Training Details drawer for AutoML models, providing:
 - Real-time progress tracking — current stage and overall completion.
 - Feature engineering insights — all features processed and evaluated.
 - Hyperparameter transparency — detailed logs of hyperparameters used.
 - Failure diagnostics — immediate access to status reports and error logs.
- **Business Impact:** Removes the black-box nature of AutoML, enabling faster debugging, informed model selection, and audit-ready explainability for production deployment.

Enhanced Conda Environment Management

- **Enhancement:** End-to-end visibility into asynchronous Conda environment lifecycles. New capabilities:
 - Unified Create status for new and imported environments.
 - YAML Preview button to inspect configuration files and execution logs in-UI.
 - Detailed failure logs and one-click Retry for failed operations without leaving the screen.
 - Validation warnings for imported YAMLS missing an explicit Python version.
- **Business Impact:** Transforms a black-box process into a transparent, self-service workflow — reducing downtime, eliminating backend intervention, and preventing silent execution failures.

Model and AI Risk Management (MARM)

- **Enhancement:** MARM introduces a structured, validator-led model validation workflow to standardize risk assessment across the organization. Features include:
 - Governed end-to-end workflow: Data Profiling & Drift Analysis, Outcome Analysis & Backtesting, Conceptual Soundness, Tiering, Findings, and Final Review.
 - Draft saving, comprehensive validation history tracking, and detailed execution reports.
 - Portfolio-level oversight via Dashboard, Model Inventory, Model 360 view, and Validation Calendar.
 - Step-level commentary and formal validation decisioning within a strictly governed lifecycle.
- **Business Impact:** Digitizes the model lifecycle to ensure all AI, ML, and Statistical models meet regulatory and internal standards — providing audit-ready executive visibility and mitigating risk associated with model failure.

Admin and Installation

Infrastructure Observability

Build-monitor components now include configurations for Prometheus, Node Exporter, and Oracle Exporter. Prometheus remains optional, giving teams the flexibility to enable advanced health monitoring based on operational needs.

Proactive Pre-Install Checks

Automated validation added for DTP placeholders, schema connectivity, service ports (DTP, UI, Jobs), and TLS-related configurations — ensuring full environment readiness before deployment.

Security and Hardening

TLS 1.3 implemented with comprehensive CVE remediation. Users can manually configure their TLS version to align with internal security policies.

Authentication and Tech Stack Modernization

Introduced LDAP-based authentication (AuthN) and upgraded to Data Studio 26.1. Java 25 is now supported, with maintained backward compatibility for Java 17 and 21.

Business Impact

Pre-install checks reduce setup time and troubleshooting overhead. TLS 1.3, LDAP support, and Java 25 keep the platform secure and compliant with modern enterprise standards, while optional Prometheus integration enables cost-effective monitoring scalability.

Limitations from Past Releases that have been Remediated in this Release

- Notebook initialization no longer fails even when the logged-in username is in lowercase.
- Graceful cleanup of data model jobs in case of abrupt shutdown of services is now handled.

Known Issues and Limitations

- Table deletion sync-up between schemas is not supported during workspace edits.
- Error log table creation fails if the column data types are **LONG**, **CLOB**, **BLOB**, **BFILE**, and **ADT** during workspace data population.
- Unable to perform the dataset cache action with the model library.
- The PDF of the model report does not contain data in the output section.
- Deployment of models that are published from the model summary screen will not promote the associated dependencies such as graphs, parameter sets, datasets, and models. However, the same works fine, if the models are published from within the pipeline UI.
- As of now, Python 3.12 does not support apache-flink completely, hence, installing Python 3.12 might display a few errors.
- Oracle-guardian-ai is no longer a mandatory library installed by mmg-python library. Instead it needs to be installed and configured in a dedicated conda environment. This is due to the limitation of the oracle-guardian-ai version compatibility and no support for Python 3.12.