



Product Introduction

金蝶Apusic分布式消息队列 V2.0.6_for_kafka

Copyright Notice

All intellectual property rights including software copyrights involved in this document have been legally registered and duly owned by Kingdee Apusic Cloud Computing Co., Ltd. Such rights are protected by Copyright Law of the People's Republic of China, Regulations on Protection of Computer Software, Regulations on Intellectual Property Protection, relevant international copyright conventions, laws, regulations and other intellectual property legislations. Unauthorized use is strictly prohibited.

Disclaimer

All copyright information contained herein is legally owned by Kingdee Apusic Cloud Computing Co., Ltd. and protected by applicable laws. The Company shall bear no liability for any information unrelated to Kingdee Apusic contained in this document. You may view, copy and print this document only within the scope permitted by Copyright Law of the People's Republic of China. No entity or individual may use, modify or republish any part or content of this document without written authorization from Kingdee Apusic Cloud Computing Co., Ltd. Any violation shall constitute infringement, and the Company reserves the right to pursue legal liabilities accordingly.

Updates to this document shall not be notified separately. You may consult the Company regarding issues related to this document. All rights not expressly granted by the Company are reserved.

Trademark Notice

 is a registered trademark filed with the National Intellectual Property Administration of PRC by Kingdee Apusic Cloud Computing Co., Ltd. Shenzhen. Exclusive rights to the registered trademark are legally owned by Kingdee Apusic and protected by law. Without written consent of Kingdee Apusic, no entity or individual shall use, copy, modify, disseminate, transcribe, or bundle and sell any part of this trademark with other products in any manner or for any reason. Kingdee Apusic shall pursue legal liabilities against any infringement of its trademark rights. All other trademarks or registered trademarks mentioned in this document belong to their respective owners.

Contents

- 1 Preface
 - 1.1 Target Audience
 - 1.2 Related Documents
 - 1.3 Technical Support
- 2 Product Overview
 - 2.1 Product Positioning
 - 2.2 Core Value
 - 2.3 Core Capabilities
- 3 Product Architecture
 - 3.1 Architecture Overview
 - 3.1.1 Console
 - 3.1.2 Core Engine
 - 3.1.3 Distributed Coordination
- 4 Product Advantages
- 5 Application Scenarios

1 Preface

This document is the product introduction for ADMQ for Kafka (Apusic Distributed Message Queue, abbreviated as ADMQ for Kafka), helping users quickly understand the features, technical characteristics, and typical application scenarios of ADMQ for Kafka.

1.1 Target Audience

This document is intended for IT informatization business leaders, R&D managers, software project managers, software architects, and operations engineers.

1.2 Related Documents

For more information about ADMQ for Kafka, please refer to the following product manual documentation suite:

No.	Manual Document	Description
1	ADMQ for Kafka Quick Start Guide	A brief introduction on how to quickly get started with ADMQ for Kafka.
2	ADMQ for Kafka Installation Guide	Detailed instructions on installing ADMQ for Kafka on various operating systems and starting/stopping ADMQ for Kafka services.
3	ADMQ for Kafka Message Engine User Guide	Detailed instructions on using, configuring, and managing ADMQ for Kafka message engine features and supporting tools.
4	ADMQ for Kafka Console User Guide	Detailed instructions on using and operating ADMQ for Kafka console features.
5	ADMQ for Kafka Development Guide	Detailed instructions on developing ADMQ for Kafka client applications in various programming languages.
6	ADMQ for Kafka Migration Guide	Detailed instructions on migrating from Kafka to ADMQ for Kafka.
7	ADMQ for Kafka Operations Guide	Detailed instructions on monitoring, operations, and security hardening for ADMQ for Kafka.
8	ADMQ for Kafka Performance Tuning Guide	Detailed instructions on performance tuning for ADMQ for Kafka.

1.3 Technical Support

ADMQ for Kafka provides comprehensive technical support services. You can obtain technical support through the following channels:

- Website: www.apusic.com
- Phone: 400-855-5800
- Email: support@apusic.com
- Kingdee Cloud Community: <https://vip.kingdee.com/?productId=73&productLineId=14&lang=zh-CN>

When seeking technical support, please provide the following information:

1. Your name
2. Company information and contact details
3. Operating system and version
4. Product version number
5. Detailed information such as logs and screenshots of anomalies or errors

2 Product Overview

2.1 Product Positioning

ADMQ for Kafka (Apusic Distributed Message Queue, abbreviated as ADMQ for Kafka) is a domestic-replacement cloud-native message middleware built specifically for Kafka users. It fully adheres to the Kafka core architecture design and feature specifications, achieving full compatibility with protocols, clients, ecosystem tools, and business scenarios. It can replace the existing Kafka cluster without modification, meeting the implementation requirements of domestic replacement projects.

2.2 Core Value

It follows Kafka's core characteristics of distribution, high throughput, and low latency, while adapting to domestic replacement software and hardware environments. It provides more stable operational guarantees, more convenient operations support, and more comprehensive security capabilities, making it the preferred solution for Kafka domestic replacement.


2.3 Core Capabilities

Core Capability	Description
Fully Kafka Protocol Compatible	Fully implements Kafka's core protocols including produce/consume protocol, partition assignment protocol, Offset management protocol, and cluster coordination protocol, with behavior identical to native Kafka.
Full-Version Client Adaptation	Supports Java, Go, Python, C++ and other official Kafka clients (2.x-3.x versions) as well as third-party open-source clients, accessible without code modification.
Seamless Ecosystem Tool Reuse	Compatible with native tools such as Kafka-console-producer/consumer, kafka-topics.sh, kafka-configs.sh, and supports data synchronization tools like Kafka Connect and MirrorMaker.
Deep Domestic Environment Adaptation	Adapts to mainstream domestic operating systems (Kylin, UOS, etc.), domestic chips (ARM, RISC-V, etc.), and is compatible with domestic databases and middleware ecosystem.
Kafka Core Feature Inheritance	Supports Kafka native core features such as topic partitions, replica mechanism, consumer groups, message compression, and Offset commit (automatic/manual).

3 Product Architecture

3.1 Architecture Overview

ADMQ for Kafka adopts the classic Kafka distributed architecture, fully adhering to the Kafka communication protocol and storage model. The overall product architecture is divided into three major modules: Console, Core Engine, and Distributed Coordination, featuring high reliability, high availability, and scalability.



3.1.1 Console

- The Console is a visual operation interface for operations personnel and administrators, providing full lifecycle management capabilities:
- Cluster Management: Connect and monitor the status of message cluster nodes.
- User Permissions: Role-Based Access Control (RBAC) to ensure operational security.
- Topic Management: Create, delete, and configure message topics, manage partitions and replicas.
- Subscription Management: Manage consumer groups, subscription relationships, and consumption offsets.
- Message Search: Query historical messages by time, key, content, and other conditions.

3.1.2 Core Engine

The Core Engine is the heart of message processing, compatible with the Kafka protocol, ensuring reliable message transmission and storage:

Kafka Protocol Access & Security Access Control

- Provides a standard Kafka protocol compatibility layer, supporting seamless migration of existing Kafka clients.
- Integrates authentication (e.g., SASL), authorization, and transport encryption (SSL/TLS) to ensure access security.

Message Processing Services

- Message Produce Service: Receives messages from producers, performs validation and routing.
- Message Consume Service: Pushes or pulls messages to consumers, supporting sequential consumption and batch consumption.
- Message Persistent Storage: Writes messages to disk, supporting high-throughput writes and efficient retrieval, ensuring no data loss.
- Metadata Service: Manages metadata such as topics, partitions, replicas, and consumer groups, ensuring cluster state consistency.

High Availability & Load Management

- Multi-Replica Synchronization: Ensures data redundancy through multi-replica mechanism (e.g., ISR), preventing single points of failure.
- Failure Recovery & Replica Migration: Automatic failover on node failure, supports replica migration within the cluster for load balancing.
- Load Balancing: Evenly distributes produce and consume requests across cluster nodes to improve overall throughput.
- Status Monitoring: Real-time monitoring of node, partition, and replica health status, providing basis for failure recovery.

3.1.3 Distributed Coordination

The Distributed Coordination module (based on Kingdee Apusic Distributed Configuration Center) is responsible for maintaining cluster consistency and state:

- Cluster Member Management: Maintains the online/offline status of all Broker nodes in the cluster.

- Metadata Persistent Storage: Stores and synchronizes core metadata such as topics, partitions, and replica assignments.
- Controller Election: Elects the Controller node in the cluster, responsible for partition management and fault handling.
- Partition Leader Election: Automatically elects a new Leader from Followers when the partition Leader node fails.
- Cluster State Change Notification: When cluster state changes (e.g., node online/offline, Leader change), sends notifications to relevant components to trigger corresponding actions.

4 Product Advantages

- Zero-Modification Domestic Replacement: Directly replace the existing Kafka cluster without modifying business code or adjusting client configurations (only need to change the connection address), with migration costs approaching zero.
- Full Ecosystem Compatibility: Kafka native tools, monitoring systems (Prometheus+Grafana), data synchronization tools, etc., can all be directly reused without additional adaptation development.
- Domestic Environment Adaptation: Fully compatible with the domestic software and hardware ecosystem, meeting government and enterprise domestic replacement project requirements, providing end-to-end domestic replacement solutions.
- Authentic Feature Set: Fully adheres to Kafka core feature specifications, supporting all familiar business scenarios (e.g., high-throughput data transmission, streaming data processing), with zero business perception during switchover.
- Enhanced Security Compliance: Adds new security features while retaining Kafka's original security capabilities, meeting domestic replacement project security compliance standards.

5 Application Scenarios

- Domestic Replacement Scenario: In government and enterprise domestic replacement projects, directly replace the existing Kafka cluster without refactoring business systems, quickly meeting domestic replacement requirements.
- High-Throughput Data Transmission Scenario: Such as log collection, business data synchronization, etc., inheriting Kafka's high-throughput characteristics to support real-time transmission of massive data.
- Streaming Data Processing Scenario: Compatible with streaming computing engines such as Flink and Spark Streaming, seamlessly connecting to existing streaming data processing pipelines.
- Distributed System Decoupling Scenario: Based on Kafka's native publish-subscribe model, enabling asynchronous communication between distributed systems and reducing system coupling.

UNIFIED SERVICE HOTLINE

4008-555-800

Kingdee Apusic Cloud Computing Co., Ltd. (referred to as “Kingdee Apusic”) was founded in 2000. Its predecessor was “Kingdee Middleware Company.” It is a basic software platform provider under Kingdee Group, a national standards-setting enterprise for cloud computing, and a core software enterprise in the national IT application innovation industry. Kingdee Apusic is an enterprise undertaking the National 863 Key R&D Program and the National Major Science and Technology Project for Core Electronic Components, High-end General-purpose Chips and Fundamental Software, and is also a basic platform provider for the national key projects of “Two Networks, One Website, Four Databases, and Twelve Golden Projects.” Its products are widely used in key industries such as government, military, finance, and energy, and it has cumulatively served more than 100,000 customers.



National standard setter for cloud computing
Basic software enterprise under Kingdee Group
Core enterprise for IT application innovation
Website: www.apusic.com

