Generated Date: 2026-10-23

Product Overview

Apusic Key-Value Database

Copyright © Kingdee Apusic Cloud Computing Co., Ltd. Shenzhen 2026. All Rights Reserved.

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

Apusic 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 Documentation
  - 1.3 Technical Support
- 2 Product Introduction
- 3 Core Features
  - 3.1 High-Performance Cache Engine
  - 3.2 High Availability Architecture
  - 3.3 Security Enhancement
  - 3.4 Operations Management
  - 3.5 Development Support
- 4 Technical Features
  - 4.1 High Performance
  - 4.2 High Reliability
  - 4.3 High Security
  - 4.4 High Compatibility
- 5 Application Scenarios
  - 5.1 Centralized Session Storage for Application Servers
  - 5.2 Hot Data Caching
  - 5.3 Distributed Lock
  - 5.4 Lightweight Message Queue
  - 5.5 Counter and Interface Rate Limiting
  - 5.6 Time-Limited Business

1 Preface

This document provides an introduction to Apusic Key-Value Database (KVDB) V2.0, helping users quickly understand the product's features, technical characteristics, and typical application scenarios.

1.1 Target Audience

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

1.2 Related Documentation

For more information about Apusic KVDB V2.0, please refer to the following Apusic KVDB V2.0 product manual documentation set:

No.	Document	Description
1	Apusic Key-Value Database V2.0 Quick Start Guide	A brief introduction on how to quickly get started with KVDB.
2	Apusic Key-Value Database V2.0 Installation Guide	Detailed introduction on how to install KVDB on various operating systems, KVDB service start/stop operations, and the product registration process.
3	Apusic Key-Value Database V2.0 Key-Value Core User Manual	Detailed introduction on how to use, configure, and manage KVDB features and supporting tools.
4	Apusic Key-Value Database V2.0 Console User Manual	Detailed instructions on using and operating KVDB console features.
5	Apusic Key-Value Database V2.0 Development Guide	Detailed instructions for developing KVDB client applications using various programming languages.
6	Apusic Key-Value Database V2.0 Migration Guide	Detailed instructions for migrating KVDB from historical versions to V2.0, as well as migrating from Redis to KVDB.
7	Apusic Key-Value Database V2.0 Operations Guide	Detailed instructions for KVDB monitoring, operations, security hardening, and other operational procedures.
8	Apusic Key-Value Database V2.0 Performance Tuning Guide	Detailed instructions for KVDB performance tuning.

1.3 Technical Support

KVDB products provide 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 requesting 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 including logs, screenshots, etc. of any exceptions or errors

2 Product Introduction

Apusic Key-Value Database Software (Apusic Key-Value Database, abbreviated as KVDB) is an enterprise-grade, high-performance key-value database software that is fully compatible with the Redis protocol and development ecosystem, enabling seamless migration without modifying existing business applications. KVDB features core capabilities including high-performance microsecond-level response, high availability with multi-replica fault tolerance, high scalability with elastic expansion, and refined web-based visual management, providing secure and reliable high-frequency data access support for critical business in high-concurrency and distributed scenarios, reducing database pressure and improving overall system throughput.

KVDB is built on Kingdee Apusic's years of middleware technology accumulation, with fully autonomous and controllable architecture from the underlying layers to upper-level functions, meeting the localization replacement requirements in information technology innovation environments, while providing enterprise-grade technical support and service guarantees.

3 Core Features

3.1 High-Performance Cache Engine

- Ultra-high throughput: Multi-threaded network I/O architecture and lock-free concurrency model, with single-node throughput up to 250,000 OPS
- Microsecond-level response: Pure in-memory operations combined with efficient data structures, achieving microsecond-level operation response
- Rich data structures: Complete support for Redis core data structures including String, Hash, List, Set, Sorted Set, Geo, HyperLogLog, Stream, etc.
- Flexible cache policies: Supports multiple memory eviction policies and key expiration mechanisms such as LRU, LFU, TTL, etc.

3.2 High Availability Architecture

- Master-slave replication: Implements data redundancy and read-write separation, improving system availability
- Sentinel mode: Provides high availability capabilities including automatic fault detection, failover, and configuration propagation
- Cluster mode: Supports data sharding storage, elastic scaling, automatic failover, meeting large-scale business requirements
- Dual persistence mechanism: Supports both RDB snapshot and AOF log persistence methods, which can be used individually or in combination, balancing data persistence efficiency and integrity

3.3 Security Enhancement

- Access control: Supports password authentication and ACL fine-grained permission control
- Transport encryption: Supports SSL/TLS mutual authentication
- Audit logging: Records complete operation logs to meet compliance requirements
- Command security: Supports renaming or disabling dangerous commands

3.4 Operations Management

- Visual management: Provides a Web console, supporting centralized and unified operations management for multi-version key-value core instances
- Cluster management: Provides a complete toolchain for cluster creation, inspection, repair, and scaling
- Performance monitoring: Real-time monitoring of key metrics such as QPS, response time, memory usage, etc.
- Slow query analysis: Automatically records and analyzes slow query commands to assist performance tuning
- Observability: Built-in Prometheus support, providing rich runtime metrics and a standard Grafana Dashboard

3.5 Development Support

- Protocol compatibility: Fully compatible with RESP2/RESP3 protocols, supports seamless integration with Redis clients
- Multi-language support: Compatible with mainstream client libraries such as Jedis, Lettuce, redis-py, go-redis, etc.
- Lua scripting: Built-in Lua interpreter, supports atomic execution of complex logic
- Complete toolchain: Provides CLI tools, benchmark tools, file repair tools, configuration conversion tools, etc.

4 Technical Features

4.1 High Performance

- Network layer multi-threaded concurrent architecture
- Single-threaded command execution ensures atomicity
- Memory-optimized data structure design
- Zero-copy technology reduces memory overhead

4.2 High Reliability

- Multi-replica data redundancy
- Decentralized cluster architecture
- Automatic fault detection and failover
- Data consistency guarantee mechanism

4.3 High Security

- ACL fine-grained permission control
- SSL/TLS encrypted transmission

4.4 High Compatibility

- Complete Redis protocol compatibility
- Mainstream programming language client support
- Standardized API interfaces
- Rich tool ecosystem

5 Application Scenarios

5.1 Centralized Session Storage for Application Servers

Provides unified Session storage capabilities for application server (such as Tomcat, AAS, etc.) cluster deployments, solving the problem of isolated Session data across cluster nodes and session state loss when users access across nodes, achieving global session sharing and unified management.

5.2 Hot Data Caching

Serves as a "front-end buffer layer" for databases, storing frequently accessed hot data in applications, reducing high-frequency query pressure on databases, and preventing database overload due to high concurrent requests. It shortens the data access path, reducing response latency from millisecond-level at the database to microsecond-level, improving overall system throughput and user experience.

5.3 Distributed Lock

Provides cross-service, cross-node mutually exclusive access control for concurrent operations on shared resources in distributed architectures. By using a distributed cache, concurrent conflicts in distributed environments can be effectively avoided, ensuring that the same resource is operated by only one service instance at a time, guaranteeing data consistency.

5.4 Lightweight Message Queue

Can be used as a message queue to implement asynchronous communication between different services. By decoupling synchronous dependencies between services, it prevents the entire link from being blocked due to a service failure. Supports peak shaving and valley filling, buffering message peaks in high-concurrency scenarios to improve system resilience and stability.

5.5 Counter and Interface Rate Limiting

Based on atomic operations, achieves precise counting of target objects (such as page views, clicks, login counts), or limits request frequency for interfaces and users. Meets real-time counting requirements in high-concurrency scenarios (ensuring accurate counting without concurrent errors); prevents interfaces from being overwhelmed by malicious requests or traffic spikes, protecting backend service resources and ensuring system availability.

5.6 Time-Limited Business

Implements various time-limited business logic through the TTL mechanism, completing operations within a specified time or automatically expiring otherwise. Suitable for scenarios such as verification codes, coupons, limited-time activities, etc.

UNIFIED SERVICE HOTLINE

4008-555-800

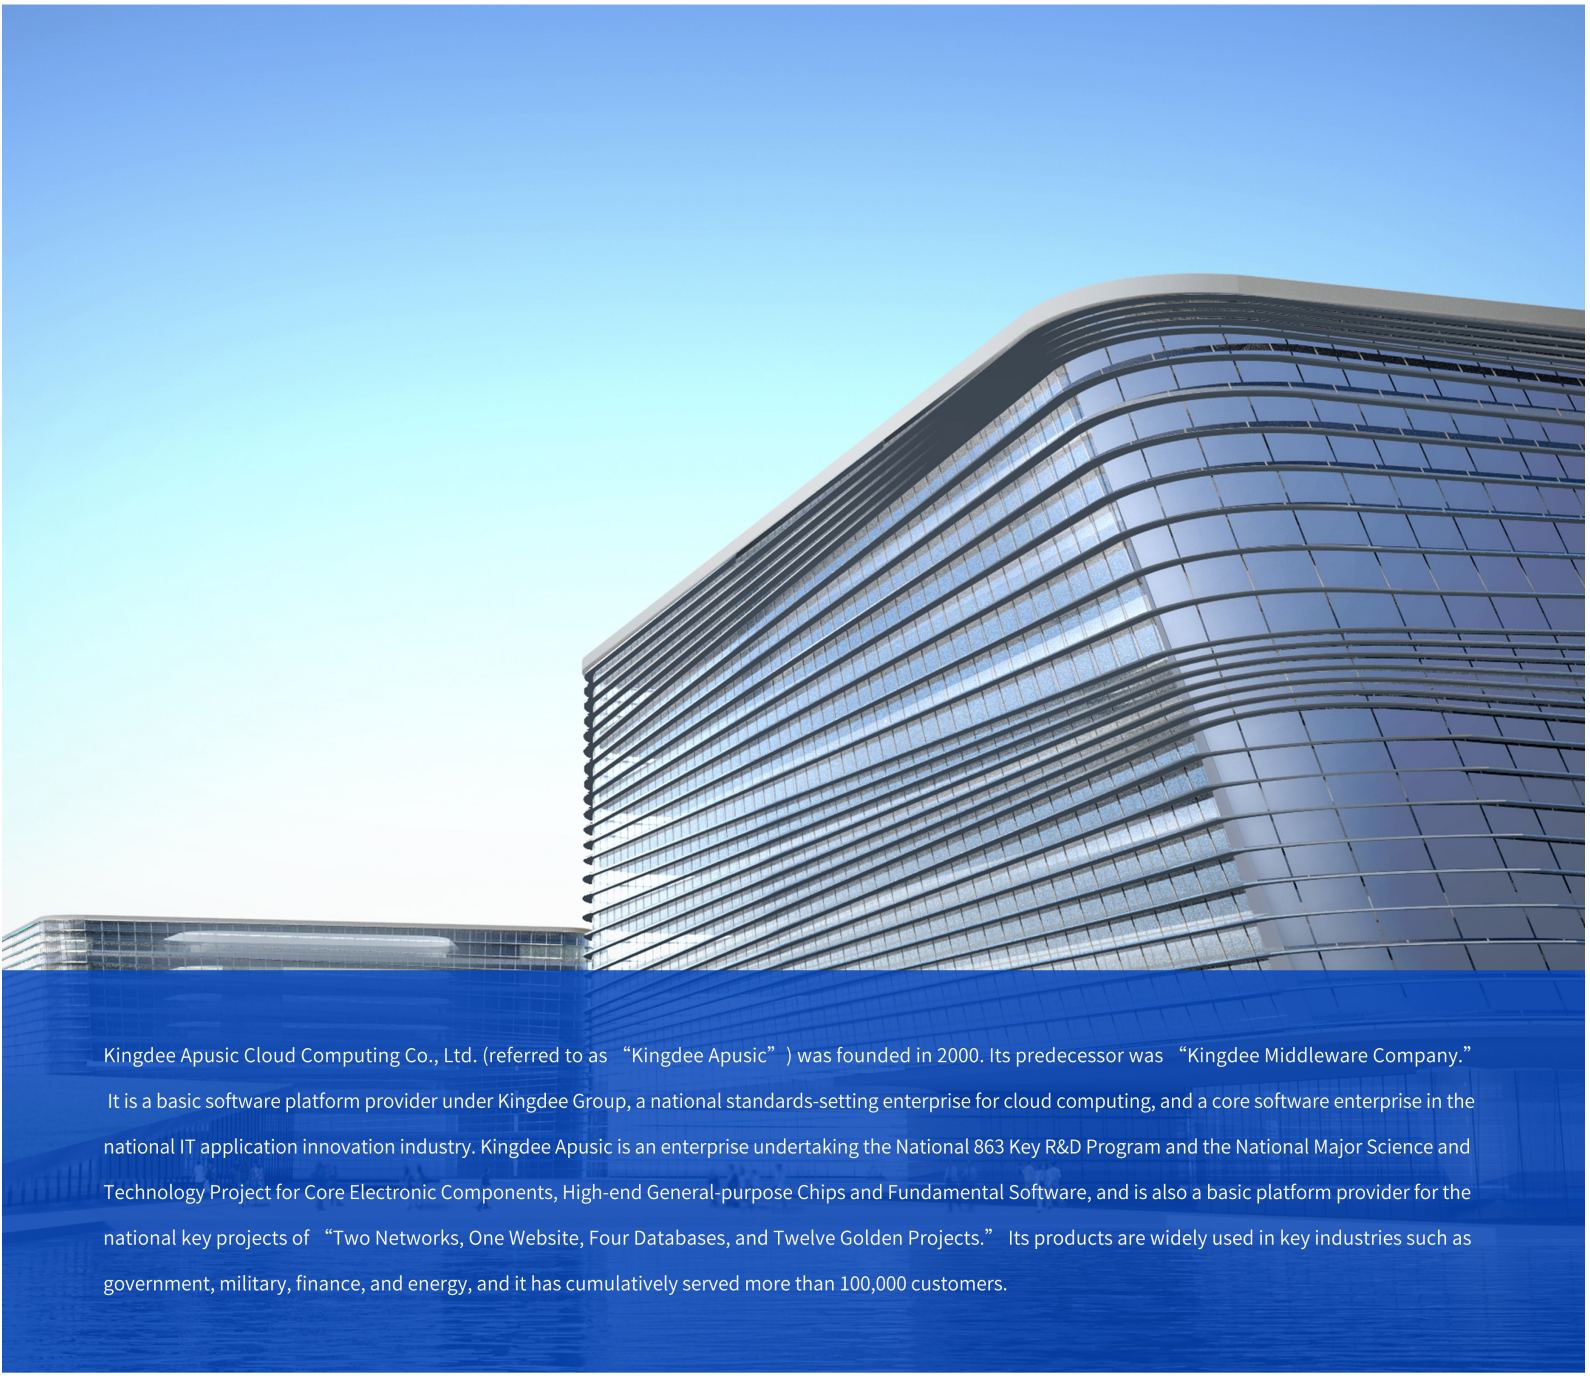

Kingdee Apusic Cloud Computing Co., Ltd. (referred to as "Kingdee Apusic") was founded in 2000. Its predecessor was "Kingdee Middleware Company." It is a basic software platform provider under Kingdee Group, a national standards-setting enterprise for cloud computing, and a core software enterprise in 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.

Apusic 金蝶天燕

National standard setter for cloud computing

Basic software enterprise under Kingdee Group

Core enterprise for IT application innovation

Website:www.apusic.com