

Introduction to Huayung CMDB

In today's complex IT environments, managing configuration items and their relationships is critical to achieving operational excellence.

Our enterprise-grade Configuration Management Database platform provides a comprehensive solution, combining industry-leading functionality with unprecedented flexibility. Designed specifically for enterprise CIOs, IT operations teams, asset management leads, SRE teams, and platform engineers, it meets the most demanding requirements of modern infrastructure management.

This platform represents a major advancement in configuration management technology, integrating six core functional pillars that work synergistically to provide end-to-end visibility, control, and intelligent support across your entire IT environment. From granular model management to AI-powered natural language interaction, each capability is designed to deliver measurable business value while reducing operational complexity.

Model Management: Laying a Solid Foundation for High Extensibility

Model management capability is the architectural cornerstone of our CMDB solution, providing enterprises with a comprehensive yet easy-to-use approach to Configuration Item (CI) organization and governance. Unlike traditional systems that require extensive customization prior to deployment, our platform comes pre-configured with over 200 CI types, enabling efficiency from day one. This rich library covers all aspects of enterprise IT assets, ranging from infrastructure components and network devices to applications, services, and business processes.

The true strength of our model management approach lies in its inheritance mechanism. By drawing on object-oriented design principles, the platform achieves dual optimization in data storage performance and model reusability. Common attributes and behaviors are defined at the parent level and automatically inherited by child CI types, ensuring consistency while eliminating data redundancy. This architecture demonstrates significant performance advantages in large-scale deployments—reducing storage overhead by up to 40% and significantly accelerating query execution speeds compared to flat model structures.

200+

Pre-built Common
CI models

40%

Storage overhead
optimization

80%

Query speed
improvement



Out-of-the-Box

Over 200 built-in standardized CI types save weeks of initial configuration costs, letting teams focus quickly on business innovation and value delivery.



Industry Best Practices

Follows proven relationship models and pre-configured topology paradigms to guarantee data accuracy and logical consistency from the initial build stage.



Advanced Inheritance Architecture

Uses an object-oriented inheritance framework to build scalable configuration models reflecting real IT structures. Attribute reuse lowers storage load while boosting both query performance and data consistency.

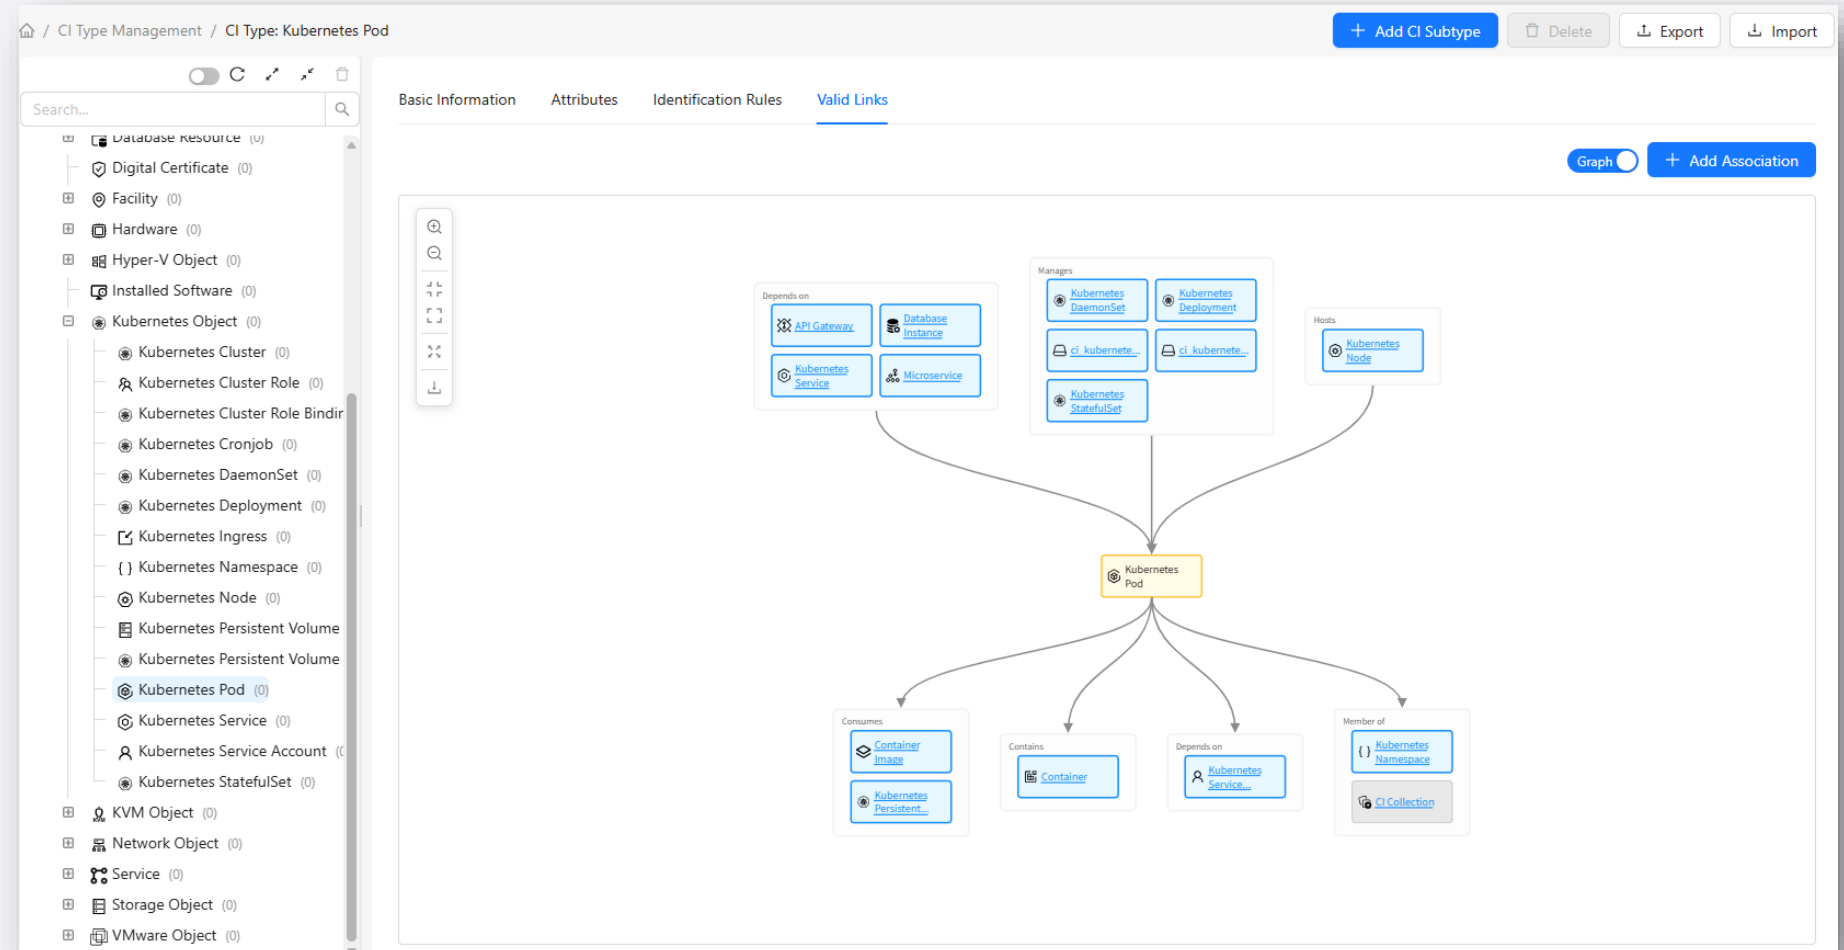


Highly Extensible

A highly flexible framework adapts to dynamic configuration needs across industries and scenarios, allowing rapid derivation of custom structures to match evolving technologies.

Model Management: Laying a Solid Foundation for High Extensibility

The platform's outstanding extensibility ensures that enterprises are never constrained by out-of-the-box models. Whether you operate in financial services, healthcare, manufacturing, or telecommunications, the system flexibly adapts to your unique operational environment. You can create custom CI types that fully support the inheritance mechanism, ensuring that extended functions maintain consistent performance characteristics and integration capabilities with native types. This flexibility proves indispensable for enterprises operating in regulated industries or possessing unique business models that require specialized configuration tracking.



Topology Construction: Effortlessly Mirroring Complex Infrastructure

Understanding the relationships and dependencies within IT infrastructure forms the foundation for effective operational management, impact analysis, and service delivery. Our topology construction capability transforms this complex challenge into an intuitive visual experience, enabling teams to map, analyze, and optimize their configuration environment with unprecedented clarity and precision.

Visual Canvas Designer

At the heart of our topology construction capability is a powerful visual canvas that allows teams to freely build topologies tailored to specific business needs. The drag-and-drop interface removes the technical barriers typically associated with relationship mapping, allowing both technical and business stakeholders to participate in topology design. The canvas supports multiple view modes—from high-level service mapping to detailed infrastructure dependency chains – ensuring the right contextual environment for different audiences and use cases.

This visual approach greatly accelerates topology creation while reducing errors. Compared to code- or form-based alternatives, team feedback indicates a 60-70% reduction in topology construction time. Immediate visual feedback helps identify logic inconsistencies, circular dependencies, and missing relationships before they can impact production systems.



Define Requirements

Identify the business services and technical components that need to be mapped



Consult Built-in Models

Select industry best-practice models matching your use cases



Customize Topology

Refine and extend the basic structure according to specific organizational details



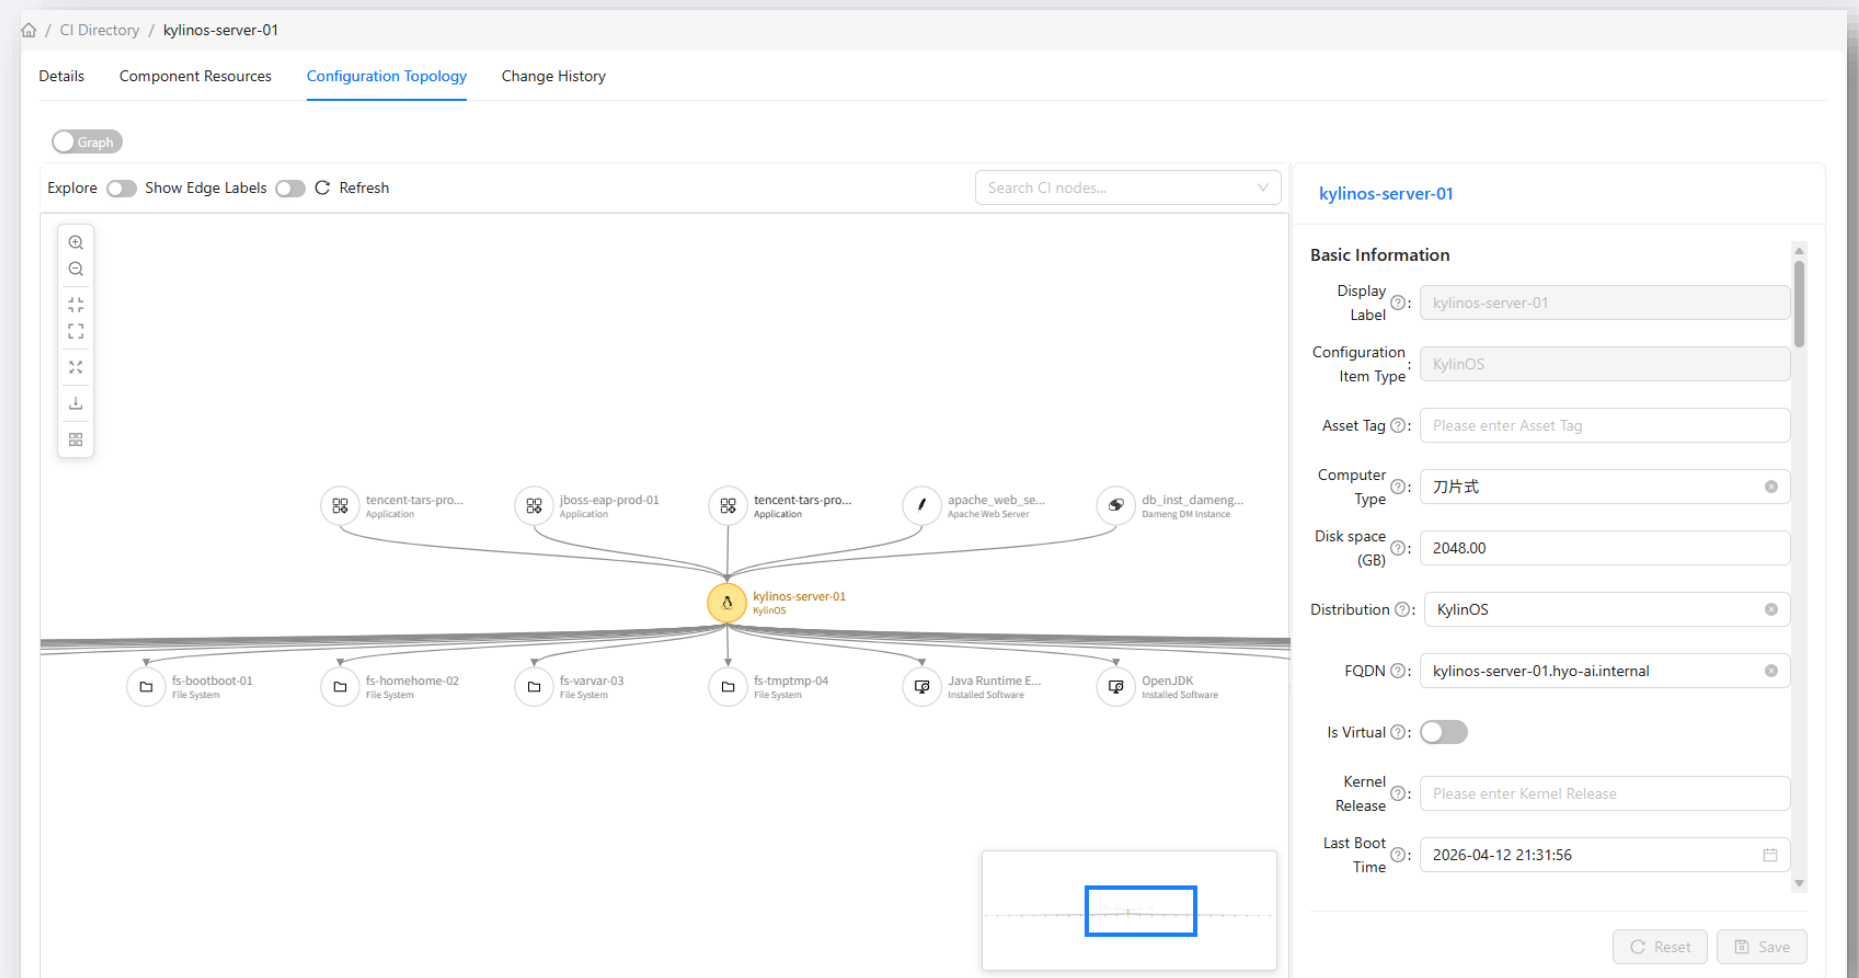
Validate & Deploy

Review relationship definitions and put the topology into active production.

Topology Construction: Effortlessly Mirroring Complex Infrastructure

Complementing the visual canvas is our built-in model reference framework, which integrates industry best practices accumulated from hundreds of enterprise-scale implementation cases. This framework comes pre-configured with recommended CI types and their relationships tailored for common scenarios – such as application service mapping, infrastructure dependency tracking, and business service modeling. By embedding these proven patterns, the reference model helps you construct topology during manual modeling that strictly align with mainstream industry frameworks like ITIL.

Furthermore, enterprises can leverage the high extensibility of these built-in models to customize CIs and relationships according to their own design standards. This not only ensures consistency across teams and projects, but also preserves valuable organizational knowledge assets.



Topology Query Language (TQL): Self-Developed Graph Query Language

Based on 20 years of industry experience, we have developed a powerful Topology Query Language (TQL) designed to unlock the full potential of your interconnected CMDB data. Unlike traditional relational database queries that struggle when handling complex relationships and multi-hop lookups, TQL is built on our proprietary graph database algorithms to provide an intuitive and efficient way to explore, analyze, and gain insights from your configuration item data.

Our syntax is optimized for the most common patterns in CMDB operations: finding upstream dependencies, tracking downstream impacts, identifying configuration drift between related items, and analyzing service topology paths. It handles multi-hop relationships, filters based on attribute values at any node along a traversal path, and aggregates results across relationship chains. Thanks to an intelligent query planner and a graph-optimized execution engine, query performance remains stable even as relationship depth increases. This flexibility fits the data consumption needs of business systems perfectly, whether for real-time operational dashboards, batch analysis processing, or integration with incident management workflows.

Intuitive Syntax

Build complex queries using easy-to-understand declarative syntax, allowing deep relationship exploration without advanced coding knowledge.

Real-Time Impact Analysis

Track dependencies in real time to quickly evaluate the downstream impacts of changes, incidents, or failures across the entire IT infrastructure.

Advanced Relationship Traversal

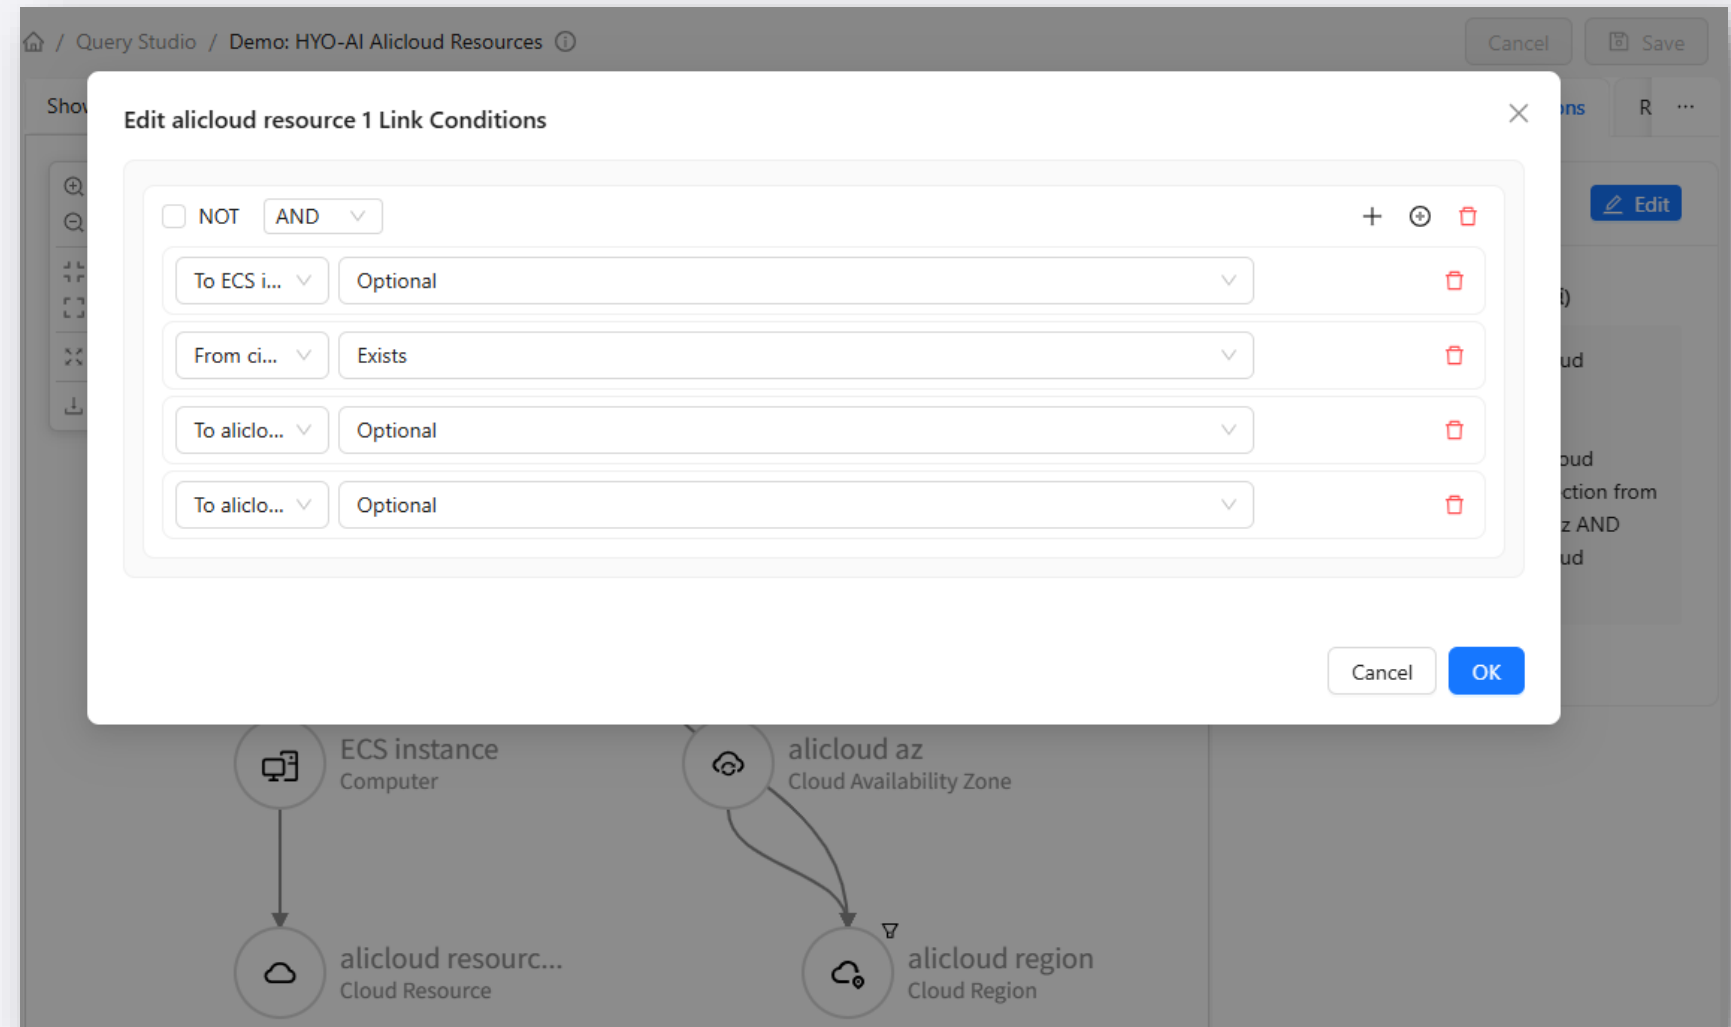
Easily explore multi-level connections between CIs to uncover complex relationship networks often missed by traditional tools.

Context-Aware Insights

Leverage the rich semantic expression of the graph model to execute smart queries that understand the meaning and context of relationships, rather than just identifying their existence.

Topology Query Language (TQL): Self-Developed Graph Query Language

By providing a dedicated graph query language, we have completely transformed how teams interact with CMDB data. It goes beyond simple data retrieval to become a dynamic analysis tool, enabling a deeper understanding of your IT ecosystem to build a more agile, robust, and optimized operational environment.



Automated Discovery: Assets and Network Topologies in Real-Time

The automated discovery engine is a fundamental component of our CMDB platform, responsible for continuously scanning and indexing your entire IT environment in real time. Through an intelligent detection mechanism combining agentless and agent-based methods, it fully identifies physical devices, virtual resources, cloud services, and edge computing nodes. Simultaneously, it accurately maps real network connections and topology relationships to build complete network topologies and asset relationship maps.

By performing real-time, precise asset inventory for all configuration items and their attributes, this engine ensures your CMDB remains a trusted single source of truth. This proactive data management style provides the underlying intelligence required for effective incident management, change control, security auditing, and capacity planning.



Automated Discovery

Uses a hybrid agentless and agent-based model to automatically detect hardware, software, network devices, and cloud environment resources across diverse IT scenarios.



Continuous Dynamic Scanning

A dual-mode update mechanism driven by scheduled tasks and events ensures the CMDB reflects the latest infrastructure state and captures environment changes instantly.



Unified Multi-Environment Coverage

Seamlessly discovers and uniformly manages various assets across hybrid clouds, on-premises data centers, IoT, and dedicated edge computing infrastructures.

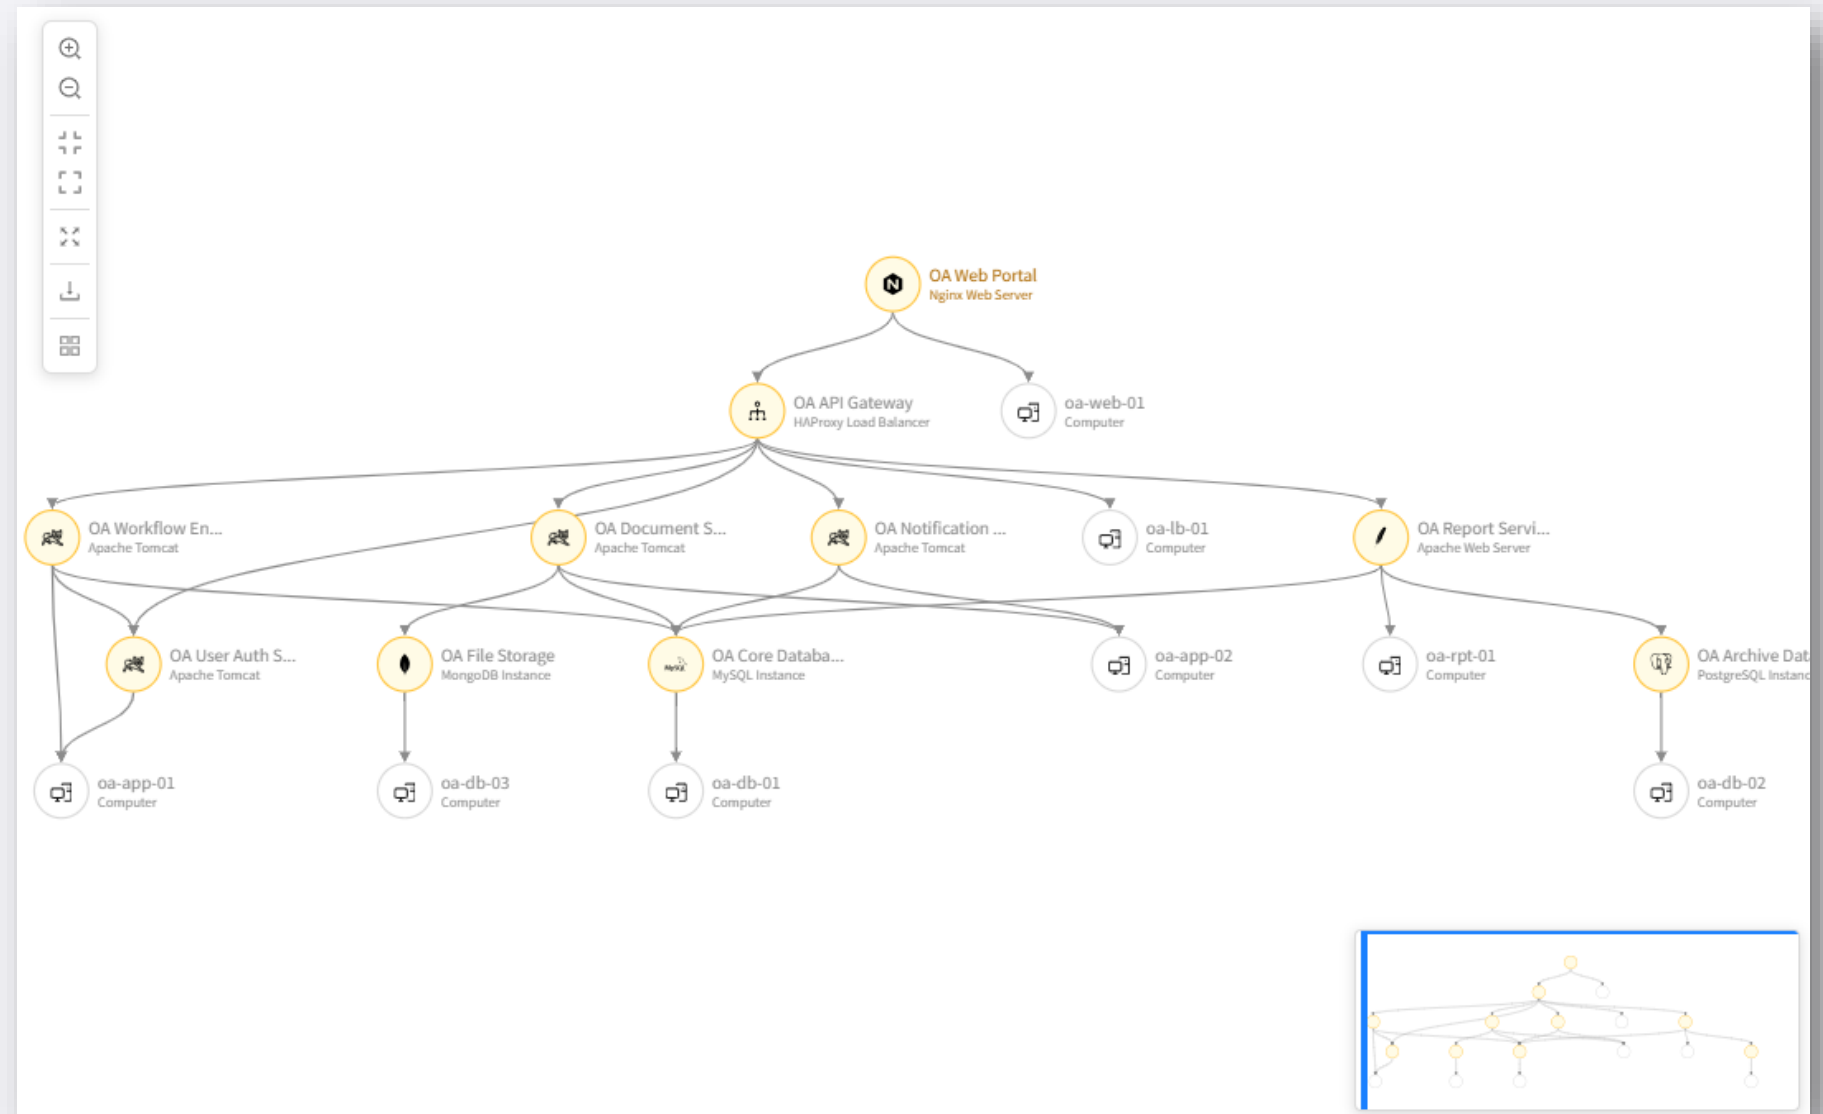


Broad Coverage

Fully supports multiple discovery protocols and data sources including SNMP, WMI, SSH, cloud platform APIs, and virtualization technologies to provide a unified data view.

Automated Discovery: Assets and Network Topologies in Real-Time

The automated discovery engine transforms traditional static asset inventories into a real-time, self-updating intelligent resource map. This provides a solid foundation for enterprise IT operations, security compliance, and digital transformation. It elevates configuration management from a passive, responsive task into a strategic advantage, offering a precise data foundation for highly efficient operations and scientific decision-making.



Intelligent Identification: The Smart Hub from Raw Data to Trusted CMDB

Whether dealing with automated discovery or integration with external systems, our identification engine acts as a critical component. It transforms raw asset data into meaningful, accurate, and actionable CIs before consolidating them into your CMDB. This advanced engine excels at blending information from disparate data sources, ensuring a unified and consistent view of your entire IT estate while overcoming the complexity of heterogeneous environments where information is scattered across different systems in multiple formats.

The identification engine strictly validates and refines data, eliminating duplicate entries and inconsistent information. By intelligently processing and correlating incoming data, it establishes a unique identity for each configuration item and precisely maps their interdependencies, laying a reliable foundation for all IT service management processes. This continuous pursuit of data accuracy ensures that your CMDB consistently reflects the evolving state of your infrastructure.

Powerful CI Identification

Uses flexible, configurable matching logic combining multiple attributes (serial numbers, IP addresses, MAC addresses, and related CIs) to achieve unique CI identification, prevent data duplication, and accurately map component relationships.

Smart Data Deduplication

Uses advanced algorithms to automatically identify and merge duplicate CI records from different sources, resolving data conflicts and unifying semantic and structural expressions for consistent, interpretable configuration data.

Authoritative Arbitration

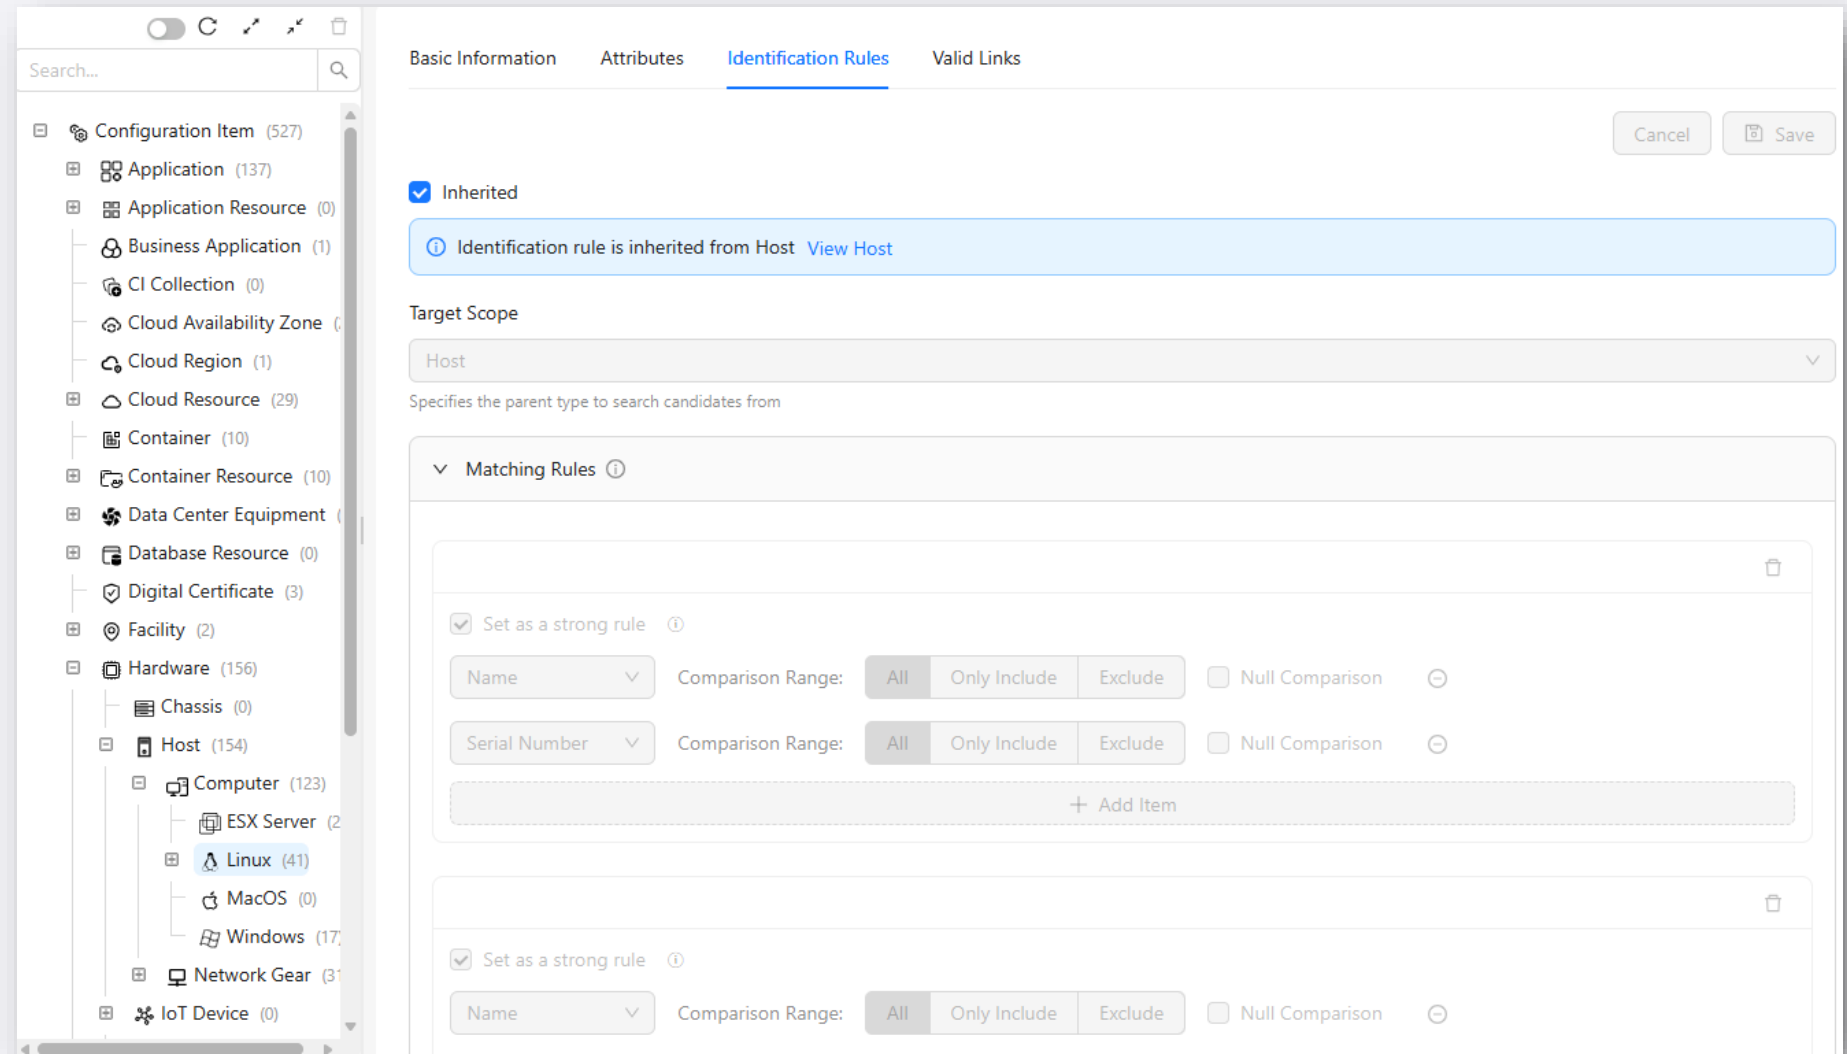
Formats and normalizes heterogeneous data across local, cloud, and network environments. Using smart arbitration algorithms and custom priority levels, it eliminates duplicates to deliver a single, authoritative configuration view.

Continuous Data Accuracy Assurance

Keeps data highly accurate and reliable through ongoing detection, ensuring the CMDB serves as a trustworthy, real-time reference baseline for all IT operations.

Intelligent Identification: The Smart Hub from Raw Data to Trusted CMDB

Serving as the critical intelligent hub connecting raw infrastructure data to a structured, trusted CMDB, the identification engine significantly improves data quality by automating complex processes like data consolidation, standardization, and unique identification. This drastically reduces the manual effort traditionally required to maintain a comprehensive CMDB. This foundational data accuracy ensures more efficient decision-making, smoother IT operations, and an overall boost in service reliability.

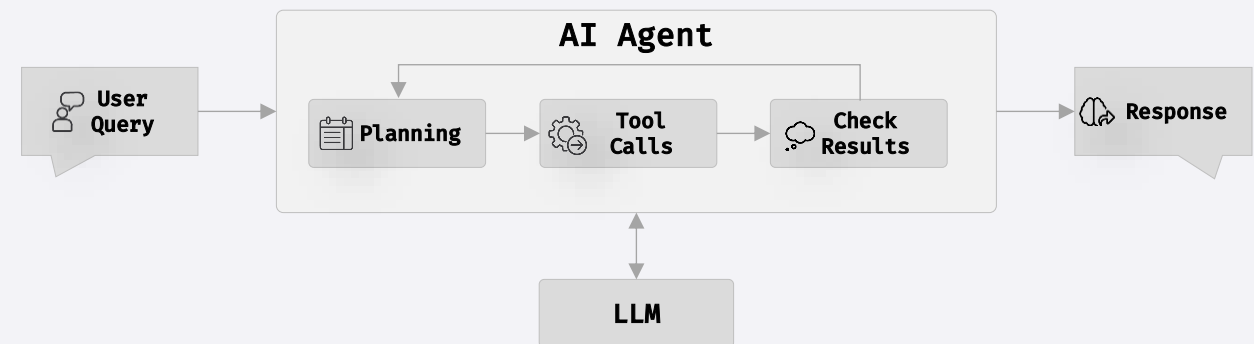


AI-Driven: Interact and Analyze Data with Built-In AI Agents

The integration of artificial intelligence marks a new era in configuration management, profoundly changing how teams interact with and extract insights from CMDB data. Our platform's built-in model context protocol integration enables seamless connection with large language models, unlocking powerful natural language interactions and flexible data analysis capabilities that traditional query interfaces could never achieve.

Users can interact with configuration data using everyday conversational language rather than formal query syntax. Questions such as "If database server DB-PROD-01 fails, which applications will be affected?" or "Which services depend on operating systems that have reached end-of-support?" can now be asked directly in natural language. The AI understands the context, user intent, and underlying data relationships to generate accurate answers.

This natural language capability significantly lowers the barrier to accessing CMDB data. Team members who lack expertise in query languages or graph traversal concepts can now independently obtain the insights they need. Business stakeholders can directly explore configuration data without relying on technical intermediaries, accelerating decision-making and fostering a data-driven culture throughout the organization.



AI-Driven: Interact and Analyze Data with Built-In AI Agents

Beyond simple query translation, AI integration empowers the system with advanced analytical capabilities. It can detect anomalies in configuration patterns, propose optimization recommendations based on relationship analysis, and even predict potential impact scenarios by reasoning through dependency chains. AI agents trained on your specific configuration data provide recommendations tailored to the unique characteristics of your environment, continuously optimizing as they learn operational patterns and outcomes.

AI interactions operate within a secure, auditable, and governance-compliant framework. All natural language queries and AI-generated insights are logged, providing full traceability for compliance objectives. Strict access controls ensure that AI-driven interfaces adhere to the exact same permission boundaries as traditional interfaces, preventing unauthorized data exposure even during conversational interactions.

The screenshot displays the 'Huayang Smart Operations Platform' interface. On the left, a navigation pane shows a tree view of configuration items, with 'Kubernetes Object' selected. The main area features a search bar and a list of configuration items. Overlaid on the right is a chat window titled 'Huayang Intelligent Operations AI Assistant'. The chat history shows a query: 'Help me query all server CI instances'. The AI response provides a detailed summary of 7 server CI instances, including their IDs, descriptions, versions, and build numbers. The response also notes that all servers are in a 'running' state and belong to a production environment. A 'View the code' button is visible at the bottom of the chat window. The input field at the bottom of the chat window contains the text 'Input issues... (Enter to send)'.

Comprehensive Benefits & Strategic Advantages

Choosing our enterprise-grade CMDB platform gives organizations far more than just configuration management software. It delivers a comprehensive solution that drives operational excellence, enables data-driven decision-making, and transforms IT operations into a strategic driver for business goals. By seamlessly blending mature functionality, innovative AI capabilities, and an extensible architecture, this platform ensures your investment meets immediate demands, delivers instant value, and provides a dependable foundation for future growth.

