Tencent Cloud

Tencent Cloud AI-Driven Cloud Migration Technology White Paper 2026

—— Empowered by MigraQ and Expert Synergy



Tencent Cloud MigraQ™ Service Center 2026

Copyright Notice

The copyright of this White Paper belongs to the Tencent Cloud MigraQ™ Service Center and is protected by law. Anyone who reproduces, excerpts, or otherwise uses the text or viewpoints of this White Paper shall attribute the source as "Tencent Cloud MigraQ™ Service Center." Violators will be held legally liable by Tencent Cloud. Except as permitted under this copyright notice, none of the contents of this white paper may be distributed, reproduced, republished or transmitted in any form or by any means without the written approval from Tencent Cloud MigraQ™ Service Center.

Legal Disclaimer

The contents of the white paper are prepared by Tencent Cloud MigraQ™ Service Center for general informational and reference purposes only. The opinions expressed in this white paper are those of the authors, and the information in this white paper may not cover all grounds and is subject to change. While Tencent Cloud MigraQ™ Service Center endeavor to keep the information up to date and correct, the information is provided “AS IS” without warranty of any kind, express or implied, about the completeness, accuracy, reliability, suitability, or availability with respect to the information, products, services, or related graphics contained in the white paper for any purpose. This white paper has not been reviewed, approved or disapproved by regulatory authorities of any jurisdiction. Any reliance you place on such information in this white paper is therefore strictly at your own risk.

YOU SHOULD CONSULT YOUR OWN TECHNICAL, LEGAL, FINANCIAL, TAX OR OTHER PROFESSIONAL ADVISOR(S) BEFORE ENGAGING IN ANY ACTIVITY IN CONNECTION WITH THE INFORMATION PROVIDED IN THE WHITE PAPER.

Contents

I. Foreword 1

II. Overview: Pain Points of Cloud Migration and the AI Opportunity 3

2.1 Pain Points and Challenges: Three Systemic Gaps 3

2.2 Bridging the Gaps: From Point Efficiency to Organizational Capability 4

III. Tencent Cloud's AI Migration Practice 6

3.1 Positioning and Mission of the Tencent Cloud MigraQ™ Service Center 6

3.2 Open Collaboration: An Efficient Service Model Built on MigraQ 9

3.3 Security: Responsibility Boundaries, the Authorization Model, and Collaborative Safety 14

3.4 Smoothness: Lifecycle AI Collaboration and Key Methods 17

3.5 Smart: The MigraQ Platform's AI-Native Architecture and Smart Engine 29

3.6 Case Studies: Benefit Assessment and Practical Evidence 33

IV. Outlook: From Project Delivery to Continuous Modernization 37

I. Foreword

Cloud migration is a systems engineering endeavor in enterprise digitalization that involves heavy investment, long delivery chains, and concentrated risk. It requires the coordinated adjustment of infrastructure, data, applications, and operations. Customers must complete the relocation of their systems while maintaining business continuity throughout the process, and must also verify data integrity, performance, security and compliance, as well as the requirements for sustained operation after moving to the cloud.

Since the release of the *Tencent Cloud Migration White Paper*, we have accumulated systematic migration processes, standards, and governance frameworks, providing customers with a clear course of action. However, as enterprise cloud adoption enters the uncharted waters of core business, the traditional approach that relies heavily on the personal experience of individual experts can no longer support scaled, high-quality delivery.

Industry research and industry-leading practices confirm this trend. Flexera's report shows that the migration bottleneck has shifted from tool supply to dependency identification and organizational coordination; **Gartner likewise points out that the skills gap has become the primary obstacle to cloud migration, and defines the fundamental transformation of service delivery as a shift from "labor-driven" to "platform-driven, AI-enabled"**; the CAICT *Cloud Computing Industry Development Research Report (2026)* also notes that cloud computing is transitioning from "resource supply" to "value delivery," with Agentic Cloud — autonomously orchestrated by AI agents — emerging as the new paradigm.

In response to this shift, industry leaders have already taken action: AWS has established the ADAPT team within ProServe to drive Agentic delivery, while IBM has consolidated its consulting assets on a unified platform. These practices demonstrate that **converting the AI opportunity into organizational capability requires a dedicated service vehicle to bear expert judgment, platform building, and governance responsibility.**

Against this backdrop, Tencent Cloud is reshaping its services for the AI era, drawing on its experience serving customers' large-scale businesses worldwide. The Tencent Cloud MigraQ™ Service Center organizes experts and collaboration resources, applies the dual-engine drive of "MigraQ + experts," and builds the MigraQ smart migration platform as the organization's AI migration capability foundation, with the goal of delivering a "**secure, smooth, and smart**" new migration experience for customers.

The core proposition of this paper is to systematically explain how the Tencent Cloud MigraQ™ Service Center transforms "individual capability" into "organizational capability", and — through the dual-engine drive of "MigraQ + experts" — to provide you with a practical migration guide for the AI era.

Keywords: AI-driven cloud migration; Agentic AI; Human-in-the-Loop; migration cutover; FinOps; MigraQ

II. Overview: Pain Points of Cloud Migration and the AI Opportunity

Drawing on migration practice and referencing research from Flexera's *State of the Cloud Report 2026*, Gartner, and CAICT, cloud migration today is not short of tools. The real contradiction is this: the efficiency gains of point solutions are being offset by systemic gaps across three layers — information, execution, and capability. AI has effectively alleviated pain points in individual stages, but converting "point efficiency" into "organizational capability" depends on systematic enablement and governance.

2.1 Pain Points and Challenges: Three Systemic Gaps

According to research, the migration bottleneck has shifted from tool supply to systemic gaps, manifested as follows:

- **Information gap (pre-migration: lack of visibility)**: Flexera notes that 54% of organizations rank "understanding application dependencies" as the top challenge; Gartner data shows 80% of organizations experience delays due to missed dependencies. Scattered assets and unclear dependencies directly lead to assessment bias, inaccurate batching, and cutover risks.
- **Execution gap (during migration: difficulty controlling execution)**: Gartner points out that over 60% of cloud project failures stem from insufficient planning or misjudged complexity. The combination of massive resource switching and scattered tools makes consistency verification and rollback decisions within the cutover window highly susceptible to control failures.
- **Capability gap (post-migration: loss of what was learned)**: Flexera data shows the "post-migration cost optimization" challenge ranking declining, reflecting cost governance being forced to "shift left." CAICT also points out that the lack of architectural experience and knowledge retention makes later optimization difficult for enterprises, and the "handover-and-leave" model leaves organizational assets unreusable.

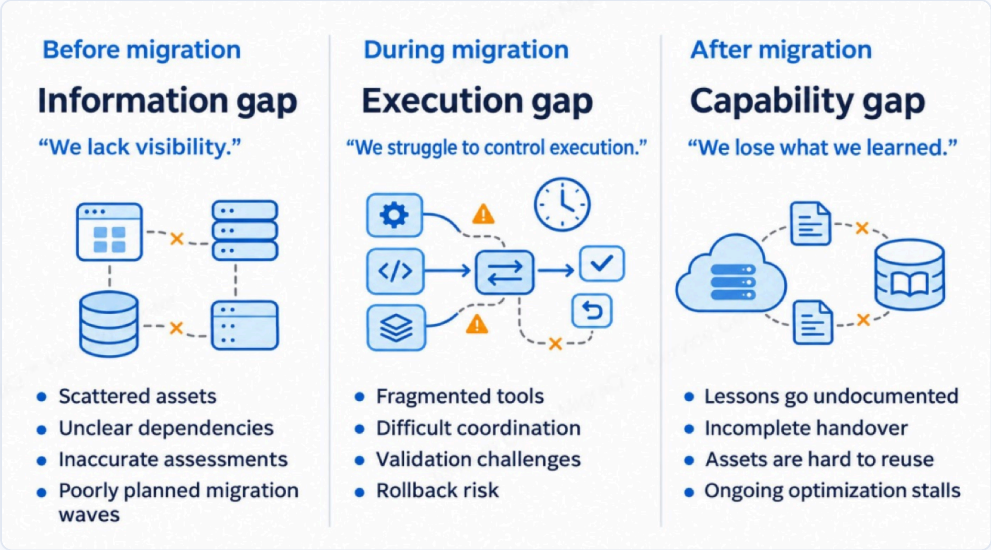

Figure 1 The three systemic gaps: information, execution, and capability

2.2 Bridging the Gaps: From Point Efficiency to Organizational Capability

With the advancement of large models and agent capabilities, AI has demonstrated clear value in bridging these gaps: at the information layer, automatically discovering network environments and extracting dependency clues of applications and resources; at the execution layer, generating verification scripts, analyzing logs, and assisting resource migration; at the capability layer, consolidating and reusing personal experience and project rules.

Yet the boundaries of AI usage are equally clear: **AI should not directly assume architectural decisions, risk acceptance, cutover authorization, or irreversible operations on production environments.** If AI output cannot enter existing approval and verification workflows, and cannot be consolidated into organizational assets within projects, point efficiency will remain at the single-team level — or even increase coordination costs by creating parallel processes. **Achieving local efficiency in a single project or team is not hard; what is hard is converting that efficiency into a stable improvement of the organization's migration service capability.**

Achieving this transformation requires four core conditions:

1. **Platform-based enablement**: Capability must not depend on individuals or single projects. Analysis paths, execution templates, verification rules, and expert judgment need to be consolidated onto a unified platform, forming a shared foundation callable across projects; it must also support customer context isolation, keeping reusable capability clearly separated from customer-sensitive data boundaries.
2. **Expert judgment and responsibility interface**: AI output requires expert review, path decision-making, and risk acceptance. Every stage of migration has confirmation points and escalation paths; AI capability must align with these nodes rather than creating a separate process. The Tencent Cloud MigraQ™ Service Center bears judgment and responsibility; the platform bears enablement and reuse — the boundary between the two must be clear.
3. **Governance mechanisms**: Including approval and audit trails, input-output contracts, responsibility boundaries, and effectiveness measurement. Converting project experience into platform assets requires clear organizational responsibility for development, maintenance, and promotion; otherwise expert time remains consumed by delivery, and platform asset building becomes a deprioritized task.
4. **Security and trustworthiness**: Migration involves customer production data and system permissions. The platform must provide environment isolation, access control, and audit logs. Configurations generated by agents should be verified in isolated environments and executed after policy-engine or manual approval; irreversible changes should not be executed directly by agents.

The above conditions show that realizing the AI opportunity is not merely a tool-selection question, but a question of how a service organization restructures its delivery model. Converting single-project efficiency into organizational capability requires a dedicated vehicle to undertake three things: **organizing expert judgment and risk responsibility, building and operating reusable platform capabilities, and establishing governance and security mechanisms to safeguard production deployment.** This is precisely the mission of the Tencent Cloud MigraQ™ Service Center.

III. Tencent Cloud's AI Migration Practice

As described above, facing the three gaps in information, execution, and capability, the efficiency gains of point tools are insufficient to cross project boundaries. Converting the AI opportunity into organizational capability requires a dedicated service vehicle to bear expert judgment, platform building, and governance responsibility. This has become an industry consensus: AWS has established the dedicated ADAPT team within ProServe to drive the agentification of its delivery system and measure business returns; IBM has consolidated consulting assets on a unified platform, emphasizing isolation and audit.

Based on its experience serving customers' large-scale business cloud migration, Tencent Cloud reshapes its delivery model through the **Tencent Cloud MigraQ™ Service Center**, driven by the dual-engine model of "**MigraQ + experts**."

3.1 Positioning and Mission of the Tencent Cloud MigraQ™ Service Center

The Tencent Cloud MigraQ™ Service Center is the core vehicle bringing together Tencent Cloud migration experts, AI tool platforms, and ecosystem resources, providing customers with end-to-end services covering assessment, design, implementation, and verification. Its mission is to join hands with ecosystem partners and, through internal and external experts and the MigraQ platform, **deliver "secure," "smooth," and "smart" cloud migration services for customers.**

3.1.1 Service Scope and 7R Path Guidance

To safeguard the quality and determinism of standardized delivery and broadly meet customer expectations, the Tencent Cloud MigraQ™ Service Center is committed to standardized cloud migration services, with its service scope focused on the infrastructure layer (IaaS) and the platform layer (PaaS). Meanwhile, referencing the industry-standard 7R classification, the implementation scope defined here concentrates on the following three:

1. **Relocate (environment-wide migration)**: Migrating virtualized environments or workloads under compatibility conditions, keeping application form as intact as possible.
2. **Rehost**: Moving existing workloads onto target resources, largely preserving the application architecture.
3. **Replatform**: Adjusting the runtime platform and necessary configurations (e.g., moving self-built databases or middleware to cloud PaaS services).

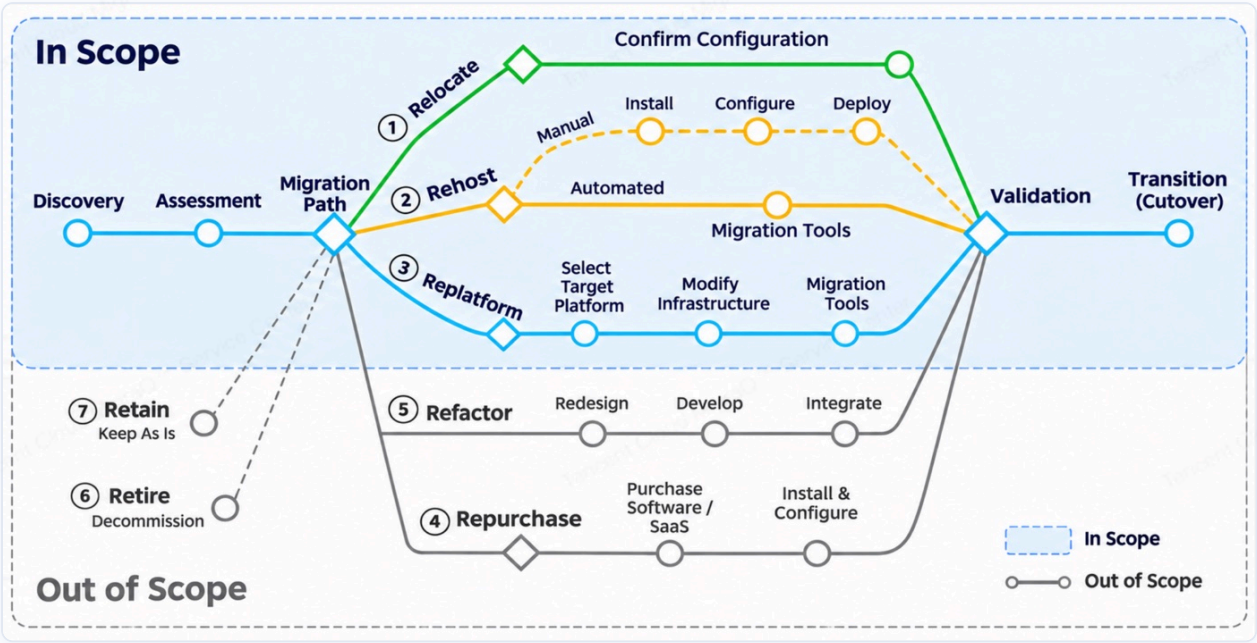

Figure 2 Service scope and 7R path guidance

To make the division of professional labor clearer, the following four strategies are outside this service's implementation scope, for the reasons given:

- **Repurchase**: Involves procuring new SaaS or commercial software to thoroughly replace existing systems.
- **Refactor**: Involves extensive refactoring of application architecture or business code, typically completed by the customer's R&D team or as an independent custom development service.
- **Retire**: An asset operation and management decision belonging to the customer. We can assist with assessment, but the final shutdown of systems is outside the implementation scope.
- **Retain**: The system is not migrated for now and keeps running in its current environment; no migration action is involved this time.

3.1.2 The Three-in-One Global Technical Operations System

Around this positioning, the Tencent Cloud MigraQ™ Service Center has built a Three-in-One technical operations system of "methodology, experts, and platform," delivering a trustworthy migration experience and assurance to customers worldwide.



Figure 3 The Three-in-One global technical operations system: methodology, experts, and platform

- **TCAF (methodology layer)**: Unified migration implementation paths, risk and delivery standards.
- **Tencent Cloud MigraQ™ Service Center (expert layer)**: Gatekeeping key strategies and risks; consolidating experience in projects.
- **MigraQ (platform layer)**: Empowering experience conversion with AI; providing replicable product support.

These three do not exist in isolation; they form a tightly closed loop during delivery:

The TCAF methodology layer provides experts and the platform with a unified "working blueprint" and compliance boundary; experts make path decisions and gatekeep risks based on the methodology standards, combined with the customer's real business constraints; the expert-verified experience and rules are then codified into the MigraQ platform layer as automatically executable standard actions (Skills). In subsequent projects, the platform lifts efficiency through automated execution while experts continue to focus on higher-order architecture design and exception handling — upgrading the entire service system from "accumulating expert man-hours" to "expert experience consolidation + AI-enabled scaled execution."

It is precisely this unified underlying operating logic that brings customers directly perceivable core value:

- **Dissolving the information gap, safeguarding business security**: Where dependency investigation once relied on experts interviewing machine by machine, AI now automatically scans multi-dimensional data and reconstructs hidden topologies, with experts reviewing only key differences. Real-world insights are quickly distilled into AI assets, so that dependencies and risks are fully identified at the solution stage — greatly reducing cutover incidents.
- **Mitigating the execution gap, enabling smooth transition**: Traditional migration was often limited by the energy and skill of individual experts. Now, verified automation tools and AI standard actions are consolidated on the platform and called across projects. AI takes on high-frequency work such as code conversion, script generation, and configuration verification, lowering the implementation threshold and error rate, and safeguarding smooth business switchover.
- **Filling the capability gap, enabling continuous smart operations**: Migration no longer resets to zero at the end of a project. Responsible owners update platform assets on a routine basis, so the handling experience of every project continuously flows back into MigraQ. Customers' post-migration continuous optimization is stably supported by AI smart inspection, realizing the evolution from "one-time delivery" to "continuous enablement."

Every project delivery not only moves the customer's business to the cloud — it also validates and iterates the platform's capability, driving migration services from personal-experience dependence toward the continuous consolidation of organizational capability.

3.2 Open Collaboration: An Efficient Service Model Built on MigraQ

Migration projects depend heavily on collaboration among business, architecture, infrastructure, and application teams. Traditional models that over-rely on on-site personnel easily lead to high on-site costs, duplicated research, and version conflicts over materials.

Based on the MigraQ platform, the Tencent Cloud MigraQ™ Service Center combines on-site support with remote expert capability, organizing customers, Tencent Cloud, and ecosystem partners to work around the same project context. Customers can flexibly choose participating teams, support forms, and service package combinations according to their technical capability, project complexity, and delivery needs.

3.2.1 The "1+N" Remote Delivery Center Model

The "1+N" remote delivery center model is a service approach in which a small number of on-site experts and a remote expert team organized on demand collaborate on the migration. "1" and "N" are generic references — actual staffing is determined by project scale, technical complexity, and the customer team's capability.

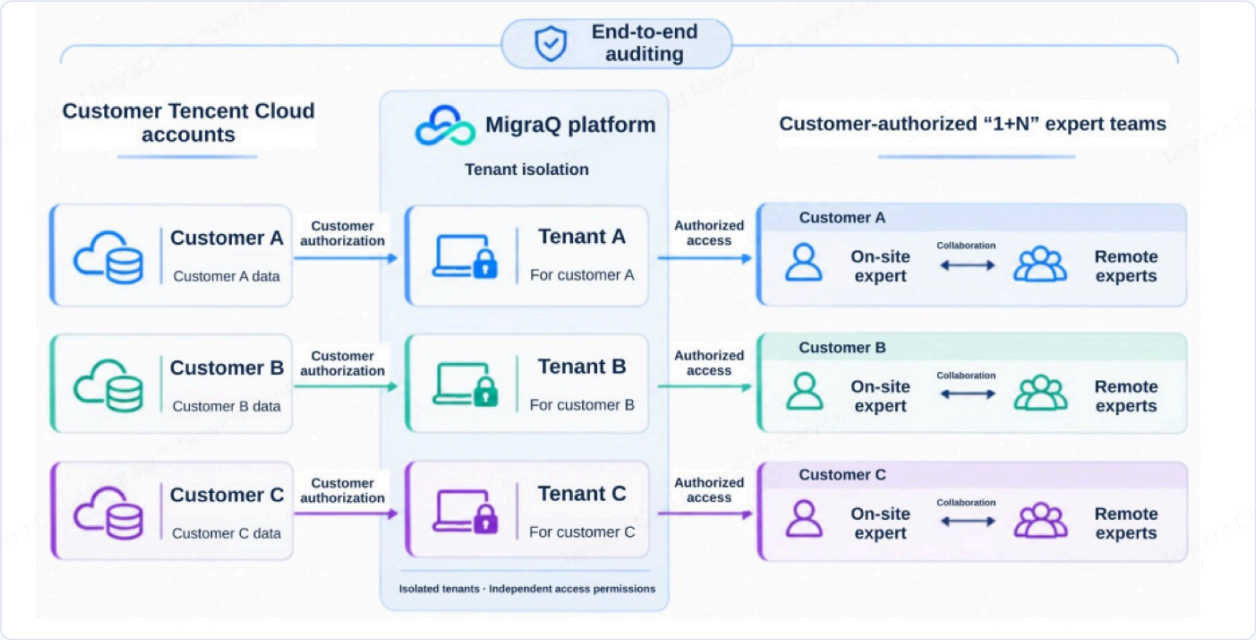

Figure 4 The "1+N" remote delivery center model

- **On-site expert (1): the context touchpoint connecting the field and the platform**. On-site experts stay close to the customer's business and operations teams, responsible for early-stage research and gathering real conditions such as business scenarios, system dependencies, network constraints, and change windows. They use MigraQ for asset sorting and preliminary analysis, and during implementation they handle on-site coordination, information confirmation, and issue feedback, ensuring the remote team works from real, complete project conditions.
- **Remote expert team (N): concentrated professional capability undertaking heavy implementation**. Remote experts rely on the MigraQ platform to carry out architecture analysis, specification mapping, migration path planning, configuration generation, and technical verification in parallel. Experts from different fields participate as tasks require, advancing work in parallel where dependencies and change windows allow.
- **Security control and asset retention**: Throughout delivery, customer data always stays in the customer's own cloud account or dedicated environment. After the customer completes authorization, experts access the MigraQ customer tenant through corresponding sub-accounts. Permissions are revoked when the project ends; all execution records and configuration assets are fully retained and handed over as the customer's dedicated digital assets.

This model controls migration risk and cost while improving delivery efficiency and data security.

3.2.2 Open System: The "MigraQ + Experts" Collaboration Ecosystem

Behind the "1+N" model, MigraQ has built a fully open expert collaboration network. Here, the definition of "expert" is broadened — whether from the Tencent Cloud team, ecosystem partners, industry experts, or the customer's own team, anyone with professional capability can become a delivery expert on the platform and participate deeply in migration co-building. Each party takes on what it does best: the customer provides business knowledge and change conditions; partners deliver the agreed localized services; Tencent Cloud provides migration methodology, product capability, and specialized technical support.

Drawing on the expert experience and AI capabilities consolidated in MigraQ, all parties can carry out migration securely and efficiently on a unified capability foundation:

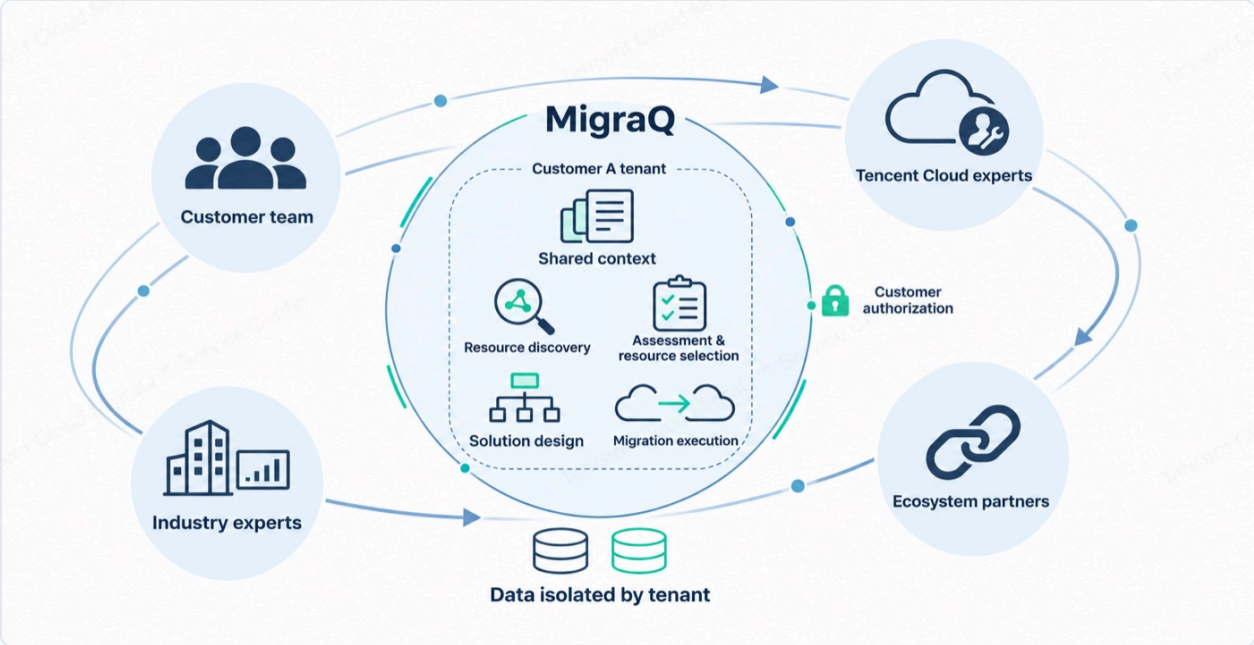

Figure 5 The "MigraQ + Experts" collaboration ecosystem

- **Seamless sharing on the same context**: When multiple teams collaborated in the past, materials were often scattered on individual local machines, easily triggering version conflicts. Now, members authorized by the customer can collaborate in real time on the asset information, business constraints, solution versions, and task results within the same MigraQ tenant instance — reducing duplicated material hand-offs and communication costs.
- **Open ecosystem alongside strict data isolation**: Open collaboration presupposes explicit data and permission boundaries. Tenants of different customers are mutually isolated; members within the same tenant receive only the permissions their roles require. A shared project context does not mean all members can access all materials or perform identical production operations. The customer holds the final say and can adjust participation scope and authorizations at any time as the project progresses.

3.2.3 Flexible Service Packages: An Open Combination of On-Demand Choices

To adapt to the open system above, the Tencent Cloud MigraQ™ Service Center provides three types of human service packages plus the platform self-service mode; customers can combine them entirely on demand according to their current capability and project complexity.

(1) Tencent Cloud Expert Support Package (senior consulting & capability supplement)

For customers who need to supplement professional judgment, this package is purchased per person-day and provides in-depth research, target architecture and Landing Zone planning, justification of major technical decisions, and tackling of complex heterogeneous environments. This package focuses on expert analysis, solution design, and technical guidance.

(2) Tencent Cloud Standard Migration Package (remote AI enablement & implementation-led)

- **Positioning**: A professional service led by Tencent Cloud experts, focused on the technical implementation and landing of the migration.
- **Core content**: Relying on the **remote delivery center** and **MigraQ's AI platform capabilities**, Tencent Cloud experts are responsible for in-depth asset discovery, migration path planning, automated script generation and execution, and **cutover-phase technical support and readiness verification**.
- **Boundary clarification**: The Tencent Cloud team's "leading role" in this package is strictly focused on migration technology and implementation paths at the IaaS/PaaS layer — it does not encompass the customer's application code refactoring. Meanwhile, during cutover Tencent Cloud provides technical assurance and support; **the final decisions on stopping business writes, switching traffic, and executing the business cutover remain with, and are executed by, the customer's business team**.

(3) Partner ASP Standard Migration Package (localized ecosystem delivery)

- **Positioning**: An ecosystem service led by a Tencent-certified ecosystem partner (ASP), with Tencent Cloud providing the platform and quality monitoring.
- **Core content**: The agreed implementation work is carried by Tencent-certified partners. The partner team works on the MigraQ platform and bears delivery responsibility; Tencent Cloud provides platform support and quality checks at key nodes as agreed.

(4) MigraQ Platform Self-Service (no human-service package)

- **Positioning**: A fully self-service mode based entirely on the MigraQ platform.
- **Core content**: Technically capable customers need not purchase human services; they use MigraQ directly to complete discovery, assessment, generation, and execution. The customer is responsible for solution review, execution authorization, and business verification.

Based on their current capability, customers can combine the above service packages into the following four typical delivery scenarios:

User Scenario	Suitable When	Service Package Combination	Who Delivers & How
Full-Stack Escort (capability supplement)	Complex architecture; need to supplement both solution design and migration implementation capability	Tencent Cloud Expert Support Package + Tencent Cloud Standard Migration Package	Tencent Cloud experts provide senior consulting and solution design; the remote delivery team carries out the agreed migration technical implementation; the customer handles business coordination, key decisions, and acceptance.
Joint Build (customer-led)	Customer has a technical team and wants to lead the project with deep involvement	Tencent Cloud Standard Migration Package	The customer leads project progress and business/application-side work; Tencent Cloud carries out the agreed infrastructure and platform-layer migration implementation and remote technical support.
Partner-Led (ecosystem delivery)	Localized service needed; prefer partner-organized implementation	Partner ASP Standard Migration Package	Partners carry out the concrete implementation; Tencent Cloud provides platform support and key-node quality checks as agreed by service terms.
Self-Service Migration (platform-empowered)	Customer capable of independent implementation and verification; wants to organize migration autonomously	No human-service package; use the MigraQ platform	The customer team completes the migration with platform capabilities and is responsible for solutions, operations, and business verification.

Once service packages and delivery scenarios are settled, both parties should further confirm the migration scope, covered phases, preconditions, deliverables, acceptance criteria, and out-of-scope work. When project scale, application changes, or business windows change, the impact on workload, cost, and schedule is assessed in the agreed manner.

All delivery modes collaborate around the authorized project context in MigraQ and are organized according to the agreed review, verification, and acceptance requirements — enabling customers to flexibly choose delivery teams while maintaining continuity of materials, process, and outcomes.

3.3 Security: Responsibility Boundaries, the Authorization Model, and Collaborative Safety

Migration involves the customer's core production data and system permissions, and more than one party participates in the collaboration. Security here rests on a model of on-demand authorization and tenant-based isolation: who participates and what they can access is decided by the customer; source-side information, authorization, and operations are all managed within a unified tenant and fully auditable end to end.

3.3.1 Responsibility Boundaries: The Customer Retains Control; Service Providers Bear Delivery Responsibility

A migration involves many participants: the customer's business and operations teams, Tencent Cloud experts, ecosystem partners, and the AI agents that execute the work. These parties collaborate in a division of labor — control over data, credentials, and key decisions belongs to the customer; solution quality and delivery responsibility belong to Tencent Cloud and partners; agents carry out execution.

- **Customer**: Retains ultimate control over data and environments. Production data stays within the customer's account boundary; credentials are kept by Tencent Cloud's Secrets Manager (SSM) — either hosted by the platform or stored in the customer's own account. Key decisions such as production changes and business traffic switching are ultimately confirmed by the customer or a designated owner.
- **Tencent Cloud**: Collaborates via the platform within authorized scope, bearing solution quality, implementation risk, and delivery responsibility; never handles customer credentials and never enters the customer environment directly.
- **Ecosystem partner**: Also collaborates via the platform within authorized scope, bearing the solution quality and implementation responsibility for what it delivers; never handles customer credentials and never enters the customer environment directly.
- **AI agents**: Execute tasks in a controlled environment; outputs take effect only after rule verification and human confirmation. Delivery responsibility remains with Tencent Cloud and partners and does not transfer because of the agent's involvement.

3.3.2 On-Demand Provisioning and Authorization: A Grant Is Access; Revocation Is a Clean Cut

The customer decides on demand who participates in delivery and which support form to enable, and authorizes accordingly.

- **On-demand provisioning**: The customer may choose Tencent Cloud delivery, bringing in ecosystem partners, or implementation by its own team. As for support forms: on-site experts, a remote expert pool, or platform-only with no human support. Enabling a party completes its authorization.
- **A grant is access**: Permissions are granted minimally per task. Credentials prefer temporary forms that expire automatically; the grant simultaneously fixes the access scope, duration, and purpose — fully recorded and auditable afterward.
- **Revocation is a clean cut**: Temporary credentials expire automatically; long-term credentials can be revoked instantly; execution environments are cleaned up automatically when tasks end. After revocation, service providers retain no access path into the customer environment.

Revocation takes effect for all participants simultaneously — the customer need not check one by one who still holds access.

3.3.3 Unified Tenant: Source Information Managed by the Tenant, Auditable End to End

All parties collaborate on the same MigraQ tenant. The migration project is carried by the tenant: source-side information from resource scans, migration tasks, deliverables, and decision records all belong to it. Source information may be collected by Tencent Cloud or partners, or provided by the customer — it is consolidated in the tenant, and all parties collaborate on the same set of assets, with files and versions retained centrally.

- **Roles combine freely**: Customer teams, Tencent Cloud experts, partners, and developers collaborate around the same tenant, combining as the project requires.
- **Authorization direction is the customer's call**: If the customer wants Tencent Cloud delivery, it authorizes Tencent Cloud; to bring in a partner, it authorizes the partner; to implement in-house, it authorizes its own team. Who does what, in what scope, is decided by the customer's authorization.
- **Auditable end to end**: Authorization changes, information access, and artifact operations within the tenant are all audited; the customer can trace by person, by time, and by operation.

When a new participant joins, only a tenant-level authorization is added — no new entry points open into the customer environment; the only boundary the customer manages is this one tenant's authorization.

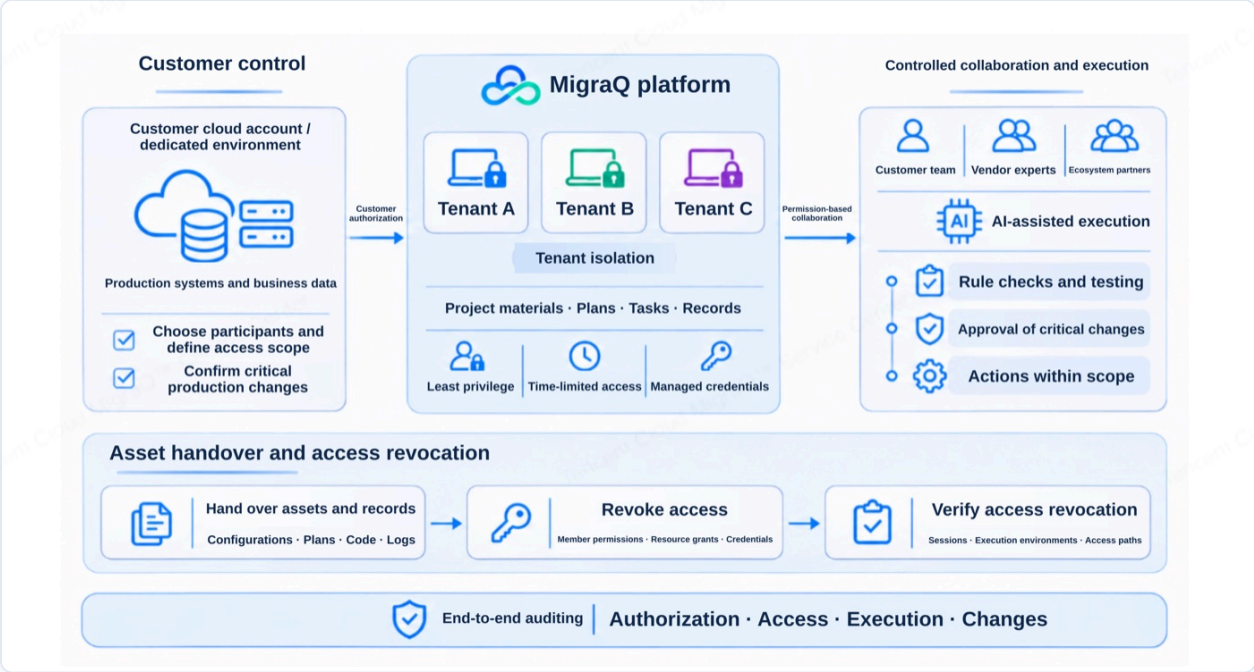

Figure 6 The collaborative security model: on-demand authorization carried by a unified tenant

3.3.4 Delivery Boundary: Asset Handover Without Risk Transfer

At project end, the asset inventory, solution and IaC code, and operation and decision records are all handed over to the customer — complete, auditable, and ready to take over; the service providers' access permissions are fully revoked at the same time. **Handing over assets does not mean handing over risk; permission revocation is the cut-off of responsibility**: the service providers retain no operations channel, and the customer enters the self-run operations stage. Subsequent cloud governance can continue on the same asset baseline consolidated during migration.

3.4 Smoothness: Lifecycle AI Collaboration and Key Methods

The "smoothness" of cloud migration means more than reducing cutover downtime: it means predictable business impact, controllable migration risk, manageable resource investment, and a migration outcome that can be smoothly handed over to the customer.

For enterprises with complex business systems and interwoven application dependencies, the real difficulty is often not moving a host or a database to the cloud, but completing the holistic migration of systems, data, networks, and organizational collaboration — without interrupting critical business.

In the *Tencent Cloud Migration White Paper 2.0*, Tencent Cloud proposed a migration implementation framework covering **research and assessment, solution design, implementation, and acceptance and summary**. That engineering logic still holds. What changes in the AI era is not bypassing necessary engineering steps, but restructuring how each step is executed: letting the platform accomplish large volumes of rule-based, repeatable work, while experts concentrate on business constraints, solution trade-offs, and key risks.



Figure 7 The end-to-end enterprise cloud migration process

In this process, AI is not an "autopilot system" that makes all decisions independently. It discovers problems, proposes solutions, and executes authorized tasks; experts are accountable for input facts, critical changes, and final business outcomes. The two collaborate through a unified asset baseline, traceable delivery records, and explicit approval nodes — balancing efficiency and reliability.

3.4.1 The Collaboration Mechanism Across the Lifecycle

To genuinely introduce AI into migration engineering, four foundational questions must first be resolved.

First, establish a continuously updated asset baseline of facts. Enterprise CMDBs, cloud platform inventories, bills, monitoring data, and manual ledgers are often inconsistent with one another. MigraQ normalizes, compares, and correlates multi-source data into a standardized asset baseline; subsequent selection, cost estimation, architecture design, and implementation verification all rest on the confirmed baseline. Business constraints and manual confirmation results should also be written back continuously, rather than relying on a single early-project scan.

Second, distinguish what machines should do from what humans must decide. Resource enumeration, multi-table comparison, configuration checks, code conversion, and document summarization suit batch execution by the platform; maintenance windows, business priorities, risk acceptance, and responsibility confirmation must be judged by business or technical owners. AI can offer comparable options, but cannot accept business risk on the customer's behalf.

Third, ground every important recommendation in evidence. Results such as specification mapping, migration strategy, and dependency inference should record data sources, applicable rules, unmet conditions, and pending confirmations. Model confidence can be used to prioritize review, but must not be equated with factual accuracy or solution pass rates. Results lacking evidence should be explicitly flagged rather than automatically entering implementation plans.

Fourth, practice controlled automation. Automatic generation does not mean automatic execution. Production writes, infrastructure changes, business cutovers, and adjustments that may affect business capacity must follow customer authorization and change management processes. The system records who approved, what was executed, what resulted, and how anomalies are handled — making critical operations auditable and traceable.

Applied to the delivery process, these mechanisms distill into five continuous phases: **asset discovery and research (Discover) → assessment and planning (Plan & Size) → environment building and application transformation (Build) → migration execution and cutover (Migrate) → acceptance, handover, and continuous optimization (Operate & Optimize)**. They are not a separate new process, but a further elaboration of the concrete work within the four-phase engineering framework.

To clearly present this "what changes and what doesn't," **we break down, across the lifecycle, the core scenarios where AI improves efficiency** — clarifying at each node the "boundary of AI efficiency" and the "expert backstop mechanism": **AI handles the "fast" parts, while people and process hold the line of "stability."**

Phase	Core Question	What AI Takes On
Discover research & inventory	What do I actually have?	Multi-source inventory reconciliation, normalization, dependency reconstruction
Plan & Size selection & cost	How to migrate, and at what cost?	Risk identification, 7R recommendation, spec mapping, TCO comparison, batch orchestration
Build construct & transform	What does the target look like?	Landing Zone / network solutions / solution drafting, IaC and code transformation
Migrate migrate & cut over	How to switch safely?	Batch script generation, workflow orchestration, anomaly analysis, cutover decision support
Operate & Optimize operate & optimize	What's still missing after migration?	Five-dimension inspection, FinOps opportunity discovery, knowledge base Q&A

3.4.2 Research and Asset Sorting: Gain Full Visibility into Migration Targets and Their Dependencies

A migration project must first answer three questions: what assets the enterprise has, which businesses these assets serve, and what the systems depend on to run properly. Resources missed or dependencies misjudged at this stage typically surface late — at testing, cutover, or even after go-live.

3.4.2.1 Multi-Source Resource Discovery and Asset Baseline Building

Enterprise resource information may be scattered across CMDBs, cloud accounts, monitoring and APM systems, financial bills, manual Excel ledgers, and historical architecture diagrams. The same instance may appear under different names, and resource configurations may contradict business ownership. Manual cross-checking is time-consuming and hard to keep current within a project cycle.

MigraQ standardizes data from different sources through resource scanning, bill and report parsing, and instance-level specification splitting, identifying duplicates, gaps, and conflicts. For example, a bill may show 29 instances of a database, while the migration plan needs per-instance version, specification, region, and dependency details. The platform links aggregated data with instance details, forming a resource view convenient for subsequent assessment.

When data conflicts, direct evidence such as instance configurations should be compared first, while other sources — bills, historical ledgers — and their collection times are retained. Irreconcilable differences go into the **asset pending-confirmation issue list** for owner verification. Only with a clear scope of assets, specifications, and ownership can subsequent estimation rest on a reliable basis.

Key deliverables: *Migration Asset Baseline Inventory*, *Resource Discovery and Asset Assessment Report*, *Asset Pending-Confirmation Issue List*. Structured JSON is the data carrier for platform correlation analysis and automated processing, and should not be conflated with the formal report for customer reading.

Common Pitfall | Treating the first scan as the final asset inventory.

The asset inventory should be a continuously updated baseline — re-verify differences after each scan round, and avoid using a one-time snapshot for all subsequent decisions.

3.4.2.2 Application Dependency and Network Topology Analysis

The resource inventory answers "what we have"; dependency analysis answers "what must move together." Enterprise applications often run across multiple business domains, databases, and network zones — missing any critical call chain can disrupt migration batch division and cutover sequencing.

Combining resource associations, configurations, and authorized network traffic observation, the platform builds application dependency relationships and network communication topologies. By analyzing call directions, protocols, communication paths, and data collection coverage, it identifies critical service relationships and potential cross-cloud access risks. Dependencies with insufficient evidence are marked as pending confirmation, rather than treating inference as fact.

Business interviews remain irreplaceable. Monitoring and traffic data reveal how systems communicate, but cannot fully answer which businesses must not stop, which time windows allow changes, or how external partners coordinate. On-site experts need to verify these business constraints with the customer and cross-validate them against system-discovered data.

Key deliverables: *Application Dependency Map*, *Network Topology and Connectivity Analysis Report*, *Business Constraints and Change Window List*. Where conditions permit, collection should cover typical business peaks, scheduled tasks, and settlement cycles — not just weekday working hours.

Practice Tip | Let traffic collection cover the full business cycle.

Collecting only weekday daytime data easily misses nightly batch processing, weekend jobs, and periodic business calls; keep a manual verification step for dependencies with insufficient evidence.

3.4.3 Migration Assessment and Strategy Formulation: Decide How, Where, and at What Cost

Once assets and dependencies are clear, migration enters the decision stage. The core here is not immediately picking target specifications, but judging which applications suit direct migration, which need transformation, and which carry technical or business risks that must be resolved first.

3.4.3.1 Compatibility and Migration Risk Assessment

MigraQ identifies lifecycle, specification-fit, address-conflict, security-policy, and cross-batch dependency risks by combining source-side OS, database, and middleware versions with target cloud product compatibility conditions and network and application dependency information.

For software versions already out of maintenance, copying to the cloud does not make the risk disappear; for resources without a fully equivalent specification, substitutes cannot be chosen merely by CPU and memory counts. The platform should clearly flag incompatible items, possible performance differences, and impact scope, and drive disposition by categories such as "blocking," "needs manual confirmation," and "monitor."

Key deliverables: *Migration Compatibility and Risk Assessment Report*, *Risk and Pending-Decision Item List*.

Practice Tip | Review evidence-poor and low-confidence risks first.

Reviews should not just browse by application name; first filter items with unclear dependencies, non-equivalent specifications, or wide impact. But confidence is only a triage aid — it cannot replace evidence and formal approval.

3.4.3.2 Migration Strategy, Target Specifications, and TCO Estimation

Migration strategy can reference the industry-standard **7R classification**, weighing options such as Rehost, Replatform, Refactor, Repurchase, Retain, Retire, or Relocate. The platform proposes candidate recommendations based on technical constraints, transformation cost, and delivery risk; experts make the final judgment combining business value and the customer's long-term technology roadmap.

Target specification mapping further resolves differences between the source side and Tencent Cloud product capabilities. Selection must compare not only specification parameters but also regional availability, service capabilities, performance headroom, and billing models. Non-equivalent recommendations must state differences and applicable conditions and be submitted for expert confirmation.

TCO estimation can only begin after the target selection is finalized. TCO should not compare cloud bills alone; it should include resource costs, migration implementation, necessary transformation, and ongoing operations per project scope, with assumptions clearly stated. Pay-as-you-go and reserved billing should be presented separately — preventing price discounts from being misread as capability equivalence.

Key deliverables: *Migration Strategy and Target Resource Mapping List*, *Migration Cost (TCO) Analysis Report*, *Migration Risk and Decision Records*.

Common Pitfall | Looking only at total TCO savings while ignoring specification downgrades.
Lower cost does not mean strictly equivalent target configurations; recommendations flagged as downgraded or substituted should be separately reviewed against performance baselines and capacity headroom.

3.4.3.3 Migration Batching and Implementation Window Planning

Large migration projects cannot simply be split evenly by server count. Batch division should account for application dependencies, stateful components such as databases and message queues, data synchronization cycles, business peaks and valleys, and customer change freezes.

MigraQ can generate candidate batches and implementation order from the dependency graph and resource attributes; experts then verify downtime tolerance, window constraints, and team concurrency with business owners. Low-confidence dependencies should be verified before deciding whether they affect batch boundaries.

Key deliverables: *Migration Batching and Implementation Plan*, *Business Cutover Windows and Responsibility Matrix*. Batch plans may be revised in later drills and reviews, but each adjustment must synchronize impact scope and owners.

Common Pitfall | Dividing migration batches simply by department, system name, or resource count.
Without verifying runtime dependencies, strongly coupled services may end up split across batches — increasing cross-cloud calls and business cutover risk.

3.4.4 Target Environment Planning and Solution Design: Build the Target Cloud Correctly

3.4.4.1 Landing Zone Planning

Moving to the cloud is not simply creating a few cloud accounts and virtual networks. How accounts are organized, who can access production resources, how business systems are isolated, how logs are retained, and how resources are brought under cost management — these seemingly basic decisions directly affect whether migration can be implemented safely and whether operations stay manageable after go-live.

A **Landing Zone** is the sustainably managed cloud infrastructure and governance framework planned and established for the enterprise before business systems move in. It typically covers account and organization structure, identity and access management, network and address planning, security baselines, log audit, and related management constraints. Think of it as the "foundation" of migration: without a unified foundation, even a successfully relocated application may be left with chaotic account permissions, blocked network access, or inconsistent governance standards.

Based on the asset baseline, target regions, and customer business constraints, MigraQ assists in generating the initial Landing Zone design and configuration list, helping architects spot conflicts in accounts, permissions, and network planning early. Cross-region, cross-domain projects must additionally consider network isolation, unified identity management, and compliance requirements. Only after customer and architect confirmation does the plan become the basis for subsequent infrastructure building.

> **Key deliverables:** *Landing Zone Plan* with its accompanying *Foundation Environment Configuration List*. The plan should clearly present the account organization, permission division, security and audit requirements, and necessary topology and design decisions; Excel and HTML are merely presentation or export formats, not the business names of deliverables.

3.4.4.2 Cloud Network Architecture and Connectivity Design

Before migration, decide how the source data centers, existing cloud environments, and Tencent Cloud interconnect; whether cross-cloud calls must persist during migration; and how network path changes are controlled at business switchover.

Combining source topology and target architecture, the platform assists in planning VPCs, CIDR blocks, routing, and access control; identifying potential address overlaps, cross-cloud bottlenecks, and security policy gaps; and informing connectivity and protection designs across dedicated lines, SD-WAN, and firewalls. Critical network policies should pass customer network and security team review, with connectivity testing where necessary.

Key deliverables: *Target Cloud Network Architecture and Connectivity Design*, *Network Address and Security Policy Configuration List*. Before generating automated deployment code, freeze the account, address, and network designs that serve as foundational constraints — avoiding rework.

Practice Tip | Confirm the Landing Zone and network design before generating IaC.

Once the account organization, address blocks, or network boundaries change, the associated Terraform modules and policy configurations may need to be regenerated and reviewed.

3.4.4.3 Migration Implementation Plan and Architecture Design

After the target-side foundation design is complete, the platform organizes migration strategy, resource mapping, dependencies, and risk-disposition requirements into an executable implementation plan. Unlike generic product manuals, a formal migration plan must correspond to specific business systems, resource objects, and migration batches.

A reviewable migration implementation plan should answer: what is the current state and target; which resources migrate; in what order; how windows are scheduled; how success is verified; when to pause or roll back; and who owns each step. MigraQ can draft migration paths and operational recommendations from the knowledge base; experts review whether content matches business constraints.

Key deliverables: *Overall Migration Implementation Plan* and, as needed, per-batch *Migration Implementation Plans*. The overall plan sets global architecture and governance; batch plans descend to concrete implementation and acceptance objects — avoiding multiple plan documents that tell different stories.

Practice Tip | Organize implementation plans by migration batch.

The overall plan describes the global view; batch plans correspond to concrete resources, review scope, cutover windows, and rollback ownership — avoiding one unwieldy runbook shared across batches.

3.4.4.4 Infrastructure as Code (IaC) and Runbooks

To reduce errors from repetitive configuration, cloud resources confirmed in planning can be created and managed using **Infrastructure as Code (IaC)**. The platform or companion development tools generate declarative configurations such as Terraform from the approved design, giving infrastructure changes version control, reviewable diffs, and repeatable deployment.

Code generation does not mean directly modifying production. The recommended sequence is "reading the confirmed baseline → generating code → conducting security and standards checks → human review → authorized execution → result verification." Changes touching production resources must additionally follow the customer's existing change management system.

Before cutover, prepare the **Application Cutover Runbook** and the **Pre-Cutover Readiness Checklist**. The former specifies every step's operations, verification methods, exception handling, and rollback paths; the latter verifies that data synchronization, network connectivity, permissions and credentials, application configuration, and owners are all ready. Rollback conditions should use measurable thresholds — not vague phrases like "roll back if necessary."

Key deliverables: *Infrastructure Automation Deployment Code and Configurations*, *Application Cutover Runbook*, *Pre-Cutover Readiness Checklist*. Code and runbooks must pass review and drills before formal implementation.

Common Pitfall | Treating model confidence as acceptance pass rate.

Confidence only helps surface candidate results needing focused review; even when the model is fairly certain, production changes must still pass rule checks, evidence verification, and authorized approval.

3.4.5 Application Transformation: Reducing Migration Friction from Heterogeneous Environments

Cloud migration is not always lift-and-shift. Systems using vendor-specific SDKs, database dialects, self-built middleware, or hard-coded addresses may need adaptation to target capabilities.

With proper authorization, development tools such as CodeBuddy assist in analyzing source code structure, interface calls, and cross-module dependencies, locating configurations, connection strings, and cloud service calls that need adjustment. For object storage, message queues, and some middleware, transformation recommendations and code drafts can be generated based on target product capabilities; database scenarios may involve compatibility analysis and conversion of table structures, stored procedures, and triggers.

Automated conversion reduces repetitive labor but cannot replace business verification. Especially for changes touching transaction calculations, complex stored procedures, or core business logic, development and database owners must confirm semantic consistency and validate results through data verification, functional regression, and performance baseline comparison.

Key deliverables: *Application Compatibility and Transformation List*, *Code and Database Transformation Outputs*, *Application Transformation and Regression Test Report*. These outputs must be linked to the corresponding applications, versions, and migration batches so issues can be traced to specific changes.

Practice Tip | Scan the code first, then verify business semantics.

Static analysis helps locate call chains, configuration items, and hard-coded addresses — but whether to delete, replace, or refactor still requires confirmation from business and development owners.

Practice Tip | Capture performance baselines before migration.

Business latency, throughput, resource utilization, and key transaction metrics should be captured pre-migration to enable reliable comparison of post-transformation performance differences.

3.4.6 Migration Execution and Business Cutover: Keep Execution Controllable and Rollback-Ready

Execution is where migration risk concentrates. Even with thorough design, live execution can still be affected by network fluctuation, data synchronization lag, configuration omissions, and shifting business windows. The migration platform's value lies not only in speed but in bringing scattered operations under a unified process.

3.4.6.1 Automated Migration and Orchestrated Coordination

Based on the approved migration plan, MigraQ invokes the appropriate migration tools or services per product and scenario, migrating hosts, containers, databases, storage, and related configurations in batches. The platform coordinates data synchronization, task execution, state tracking, and result verification, recording task inputs, executors, timing, status, and exceptions.

Repeatable, recoverable steps may be automated within authorized boundaries; production switching, sensitive configuration changes, and other high-risk operations must pass explicit approval and confirmation. This preserves AI's efficiency advantage while keeping human control at critical nodes.

Practice Tip | Review diffs before authorizing batch changes.

Terraform or scripts should first produce a plan and change diff for review — never "generate and execute"; archive versions, execution records, and recovery evidence along the way.

Practice Tip | Sequence migration by dependency.

Similar resource counts do not mean similar business risk; prioritize aligning call relationships, data synchronization, and business window coordination.

3.4.6.2 Exception Response, Cutover, and Rollback

Before the cutover window opens, experts run drills per the *Application Cutover Runbook*, check the *Pre-Cutover Readiness Checklist*, and confirm the approver, executors, and rollback owner. During execution, the platform correlates monitoring and execution logs to specific steps, providing diagnostic clues and handling suggestions when anomalies arise.

Whether to proceed, pause, or roll back should be judged by pre-agreed metrics and responsibility mechanisms. For example, when the error rate of critical business interfaces, database sync lag, or data consistency checks breach thresholds, the corresponding pause or rollback process should trigger — not continue merely on "it should recover soon."

Key deliverables: *Migration Execution and Change Records*, *Cutover Drill and Readiness Assessment Records*, *Migration Verification and Exception Handling Report*. Together these form the audit evidence of migration implementation.

Common Pitfall | Treating the formal cutover window as an ad-hoc debugging window.

Known anomalies, pause conditions, recovery steps, and escalation paths should be settled in the runbook and drills; handling unforeseen issues on the spot still must follow risk escalation and authorization mechanisms.

3.4.7 Acceptance, Asset Handover, and Continuous Optimization: Make Migration Outcomes Last

Migration completion does not mean delivery is over. Customers still need to confirm business availability, data integrity, and performance targets — and be able to run daily operations independently in the new cloud environment.

The acceptance stage turns the goals and baselines set before migration into verifiable checks, focusing on business functionality, data consistency, network connectivity, performance, security configuration, and necessary recovery capability. Discovered issues need clear owners and closure criteria — "task executed successfully" is no substitute for business acceptance.

Unlike traditional delivery, where experience stays on personal laptops or in chat logs, MigraQ continuously archives the confirmed asset baseline, target architecture, IaC code, change records, runtime configurations, and runbooks into digital assets the customer can receive, understand, and maintain. Handover is not merely file transfer: asset versions, management permissions, operational responsibilities, and future update mechanisms must all be clarified.

Key deliverables: *Migration Acceptance Report*, *Cloud Operations and Handover Package*, *Outstanding Issues and Continuous Optimization Recommendations*. The handover package includes the asset inventory, architecture materials, configurations and IaC files, runbooks, and necessary implementation and audit records; sensitive credentials must be managed separately per the customer's security policies — not delivered as ordinary documents.

After migration, and with customer authorization, cloud governance can continue on the same asset baseline — including configuration checks, security risk identification, performance optimization, and cost analysis — turning a one-time migration project into sustainable cloud operations.

3.4.8 Summary: Stable Engineering Standards, Continuously Upgrading Work Methods

AI has not changed the fundamentals that migration engineering requires rigorous assessment, thorough design, controlled implementation, and strict acceptance — but it has changed how this work is done. Multi-source asset data is normalized and compared faster; architecture and migration plans can be drafted quickly on established standards; code and execution tasks are orchestrated more efficiently; experts, in turn, devote more energy to complex dependencies, business constraints, and key risk judgment.

Migration delivery thus evolves from serial work dependent on a few experts' personal experience toward a collaboration model of **unified asset baseline, AI-assisted analysis, expert decision gatekeeping, platform-controlled execution, and continuous consolidation of outcomes**. What is finally handed to the customer is not just a business system running in the cloud, but a traceable, maintainable, and continuously optimizable cloud asset foundation.

3.5 Smart: The MigraQ Platform's AI-Native Architecture and Smart Engine

Sections 3.3 and 3.4 established the security constraints and working methods of migration engineering; these requirements must ultimately be carried by the platform through engineering means. MigraQ is the capability foundation of the Tencent Cloud MigraQ™ Service Center, having evolved from a page-driven migration tool into an agent-centric, end-to-end migration delivery platform.

Its core design philosophy: convert the industry know-how of Tencent Cloud's top experts into code-based rules and migration-domain skills at the platform's foundation. MigraQ understands requirements and orchestrates skills; computation and execution are delegated to deterministic tools and data. The model understands and orchestrates — the basis for every conclusion comes from tools and data.



Figure 8 The MigraQ platform's AI-Native layered architecture

3.5.1 Platform Architecture: Agent Engine and Execution Tools

MigraQ adopts a layered architecture: the access layer provides a unified entry point for different roles; the agent engine handles understanding and orchestration; the execution tool layer carries the concrete analysis and migration capabilities; and migration clusters execute within the customer's environment.

- **Agent engine**: Runtime environments and sessions are isolated per customer; skills and tools load on demand. Layered memory separates sessions, projects, and reusable knowledge — the asset inventory, plans, and decisions of the same project persist across collaboration sessions.
- **Execution tools**: 50+ skills covering the full migration lifecycle; capabilities such as pricing queries, specification recommendation, solution design, and migration implementation are provided by dedicated tool services.
- **Separation of control plane and data plane**: The agent plans and orchestrates; migration clusters deployed in the customer's Tencent Cloud intranet execute for real.

3.5.2 Coverage: Four Workflows Across the Pre- and Mid-Migration Stages

MigraQ supports 8 major public clouds and IDC environments as migration sources and can identify and benchmark 240+ cloud products. Platform capabilities are organized into four workflows, corresponding to the Discover-to-Migrate phases in 3.4; the post-migration Operate & Optimize phase is carried by CloudQ (see 3.5.4).

- **Resource Discovery**: Collect via cloud account authorization, or import bills and resource inventories; normalize multi-source data, sort out resource associations and dependencies, and form the asset baseline.
- **Mapping & Costing**: Benchmark Tencent Cloud products by specification, regional availability, and billing model; complete source-to-target TCO estimation based on real-time pricing.
- **Solution Design**: Generate the Landing Zone, network plans, migration batches, and cutover-rollback plans; output reviewable architecture views and IaC code.
- **Migration Implementation**: Schedule migration tools to migrate hosts, containers, databases, storage, and configurations such as gateways and load balancers — covering 17 resource types; dangerous operations require double confirmation, and traffic switching requires human confirmation.

3.5.3 The Smart Core: Codified Experience and Anti-Hallucination

Enterprise migration demands high determinism of conclusions — selection, pricing, and risk judgment cannot rely on the model's open-ended generation. MigraQ's approach narrows the model's room for improvisation, grounding key conclusions in real-time data and explicit rules:

- **Real-time key data**: Product specifications, available inventory, and prices are queried in real time via tools; resource information follows scan results. Results must pass consistency checks before being written into the asset baseline — anything that cannot be matched to real resources is rejected.
- **Rules codified in skills**: For specification recommendation, the platform prefers the newest generation of instances at similar prices, and supports strategies such as performance-first, cost-first, and strict benchmarking. Cost estimation and risk determination are likewise performed by deterministic logic within skills.
- **Clarify first, then execute**: Users describe needs in natural language (e.g., "estimate the cost of moving these machines to the cloud"); after intent recognition, the agent seeks confirmation only on key inputs that affect results, proceeding on stated default assumptions only where necessary.
- **Reviewable output**: Key recommendations carry data sources, applicable rules, and degradation or substitution notes for architect review; items with insufficient evidence are explicitly flagged and excluded from implementation plans.
- **Continuously updated rules**: Anomaly handling and expert review conclusions from projects are consolidated into new skills and rules through a "develop → upload → verify → integrate" co-building process, ready for direct reuse in subsequent projects.

3.5.4 Tool Collaboration: From Code Transformation to Cloud Governance

Beyond resource migration, applications often need transformation beforehand, governance continues afterward, and complex projects require multiple experts. MigraQ works with other Tencent Cloud AI products to cover these stages:

- **CodeBuddy**: Assists code understanding, hard-coded cleanup, and cross-cloud SDK adaptation during transformation; generated code proceeds in the order "read → generate → policy check → human approval → execute."
- **CloudQ**: After migration, reuses the asset baseline to continuously inspect security, reliability, performance, cost, and service limits — supporting FinOps and architecture governance.
- **WorkBuddy**: Serves as the expert collaboration entry, where multiple migration expert agents collaborate on demand to support task orchestration for complex projects.

3.5.5 Real Results: Key Phases Compressed from Days to Hours

The value of platform capability ultimately shows in how delivery phase durations change. Based on the Tencent Cloud MigraQ™ Service Center's project statistics, the main phase comparisons are:

Phase	Traditional Approach	With MigraQ
Resource information collection	2-3 days	0.5 hours
Specification benchmarking & selection	1-2 days	1 hour
Migration plan generation	2-3 days	2-3 hours
Configuration-type migration	5-6 days	0.5 hours
API gateway configuration migration	7 days	1 hour

In benchmark testing, the processing time of a single migration task dropped from about 240 minutes manually to 25 minutes. Note that phase-level speedups do not shrink project cycles proportionally — cutover windows, business verification, and customer decisions still follow the project's own rhythm. In one e-commerce customer's full migration project, the overall cycle was shortened from 50 days to 32 days.

Summary MigraQ carries the full migration process with a unified entry point and context, covers the pre- and mid-migration stages with four workflows, constrains model output with real-time data and skill rules, and extends capabilities to application transformation and post-migration governance through CodeBuddy and CloudQ. It reduces delivery's dependence on the personal experience of a few experts, enabling complex projects to be replicated under a unified standard.

3.6 Case Studies: Benefit Assessment and Practical Evidence

This section presents six real-world cases demonstrating the practical value of the "MigraQ + Experts" model. Ranging from large-scale multi-cloud consolidation to high-frequency specialized scenarios, these cases illustrate how AI-driven migration eliminates blind spots, ensures controlled execution, and delivers measurable gains in efficiency, cost, and risk mitigation.

Case 1: Southeast Asian Telecom's Ultra-Large-Scale Multi-Cloud Consolidation

Project background: The customer needed to consolidate complex workloads distributed across GCP, AWS, Huawei Cloud, and local IDCs onto Tencent Cloud — involving 1,200 microservices, 1,100 APIs, and 15TB+ of core data. The core requirements were business stability, zero user complaints, and an extremely tight delivery timeline.

AI's work: **Cross-cloud asset discovery and dependency mapping** — MigraQ automatically scanned the heterogeneous cloud environments (GCP/AWS/Huawei/IDC), reconciled multi-source asset ledgers, and reconstructed the complex cross-cloud call dependencies from traffic and configuration data to generate topology graphs; **Wave planning and code conversion** — based on the dependency graph, AI automatically generated the optimal migration wave plan, and CodeBuddy performed automated code conversion and adaptation for SDK call and script interface differences across cloud vendors; **Execution assistance and verification** — during migration execution, WorkBuddy assisted task orchestration while MigraQ automatically verified configuration consistency; **Experience distillation** — the conversion rules and check items for common multi-cloud environments were consolidated into 20+ dedicated migration Skills during the project.

Delivery outcomes: The inventory analysis and assessment that traditionally took 4 weeks was completed in 3 days with AI assistance; code and script conversion time was reduced by 70%. The project was delivered in 4.5 months (vs. an industry average of ~9 months), with 0 complaints and 0 rollbacks, saving the customer 50% of cloud costs.

Case 2: Smooth Cloud Migration of Complex Heterogeneous IDC with an L2GW Breakthrough

Project background: The customer needed to move local IDC workloads spanning multiple data centers — mixing physical machines, VMs, containers, and heterogeneous operating systems — to the cloud. To preserve the stability of legacy systems, the customer required a Layer 2 network extension bridging the physical boundary between the IDC and the cloud, keeping IP addresses unchanged during migration to avoid architecture refactoring risks.

AI's work: **Agentless asset discovery and topology restoration** — without installing any agent, MigraQ automatically scanned the hybrid operating forms across multiple IDCs and precisely identified hidden application dependencies through traffic analysis; **Optimal wave and configuration generation** — AI generated the batching plan based on the dependency graph and assisted in generating the configuration parameters required for the underlying Layer 2 network (L2GW) interconnection and host migration; **Execution-phase state verification** — during migration, source and target system states and network connectivity were automatically compared and verified.

Delivery outcomes: Configuration-type operations shrank from weeks to minutes, overall migration execution efficiency improved by more than 50%, achieving zero business interruption and a seamless, imperceptible cutover.

Case 3: A Leading Consumer Electronics Company (Rapid Solution and Architecture Generation)

Project background: After completing the source resource inventory, the customer needed to design the overall cloud architecture diagram and migration plan. Traditionally, this required 1-2 days of manual offline work by architects.

AI's work: Experts imported the resource inventory into MigraQ, where the platform's built-in architecture graph and rules engine automatically identified resource dependencies and generated a visualized cloud architecture diagram and a complete draft migration plan in one click.

Delivery outcomes: The solution design cycle was dramatically shortened from 1-2 days to within 1 hour; experts only needed to review core parameters and make business-level adjustments, greatly accelerating project kickoff.

Case 4: A Cross-Border E-commerce SaaS Provider (Mass Migration of K8s Objects)

Project background: The customer needed to migrate an Alibaba Cloud ACK cluster to Tencent Cloud TKE in its entirety, involving tens of thousands of K8s objects (Deployments, Services, etc.) and thousands of images to be moved and adapted to the Tencent Cloud environment. Manual operation was extremely tedious and error-prone — traditionally, it was almost impossible to move by hand.

AI's work: The platform parsed the source K8s resource manifests through automated scripts, automatically executed image synchronization and configuration adaptation, and completed the full migration of cluster objects in one click.

Delivery outcomes: Tens of thousands of K8s objects and thousands of images achieved a fully automatic, smooth transition, completely eliminating the heavy manual effort of one-by-one comparison and reconstruction, and ensuring seamless migration at the container layer.

Case 5: An IoT Enterprise (Specialized Migration of the Cloud-Native Gateway)

Project background: Migrating from Alibaba Cloud ALB to Tencent Cloud's cloud-native gateway involved extensive configuration work: forwarding rule sorting, service route creation, rule format conversion, SSL certificate synchronization, rewrite plugin configuration, and automatic binding with TKE services. Manual per-item configuration took about 1 week.

AI's work: The platform launched specialized migration capabilities for Layer 7 network components, automatically reading the source ALB configuration, converting it through the rules engine, and syncing it to the Tencent Cloud gateway — achieving one-click migration and automatic binding of configuration items.

Delivery outcomes: What once required 1 week of manual configuration was completed in 10 minutes with one click, greatly reducing personnel effort and avoiding the syntax errors and omissions that manual configuration so easily introduces.

Case 6: A Cross-Border Retail Group (Rapid Multi-Cloud Account Asset Inventory)

Project background: The customer had a complex multi-cloud environment (3 Alibaba Cloud accounts, 1 Volcano Engine account, and 1 UCloud account). Traditional manual resource research, ledger sorting, and reconciliation across cloud accounts typically took 2-3 days.

AI's work: Through MigraQ's multi-cloud connectors, all 5 accounts were authorized and connected; the platform automatically performed full resource scanning, field standardization, and multi-source ledger reconciliation.

Delivery outcomes: Asset inventory work was shortened from 2-3 days to half a day (one morning), helping the customer quickly map its asset landscape and providing a highly timely and accurate data foundation for subsequent cloud migration planning.

Summary The cases above collectively demonstrate that the core value of the AI-plus-platform combination lies not in replacing experts, but in bringing more complete information into decisions (eliminating blind spots), keeping execution deterministically controlled (zero interruption / high efficiency), and steadily distilling experience (Skill reuse). Delivery cycle time, failure rate, and asset reuse rate are the core metrics for measuring the effectiveness of this mechanism.

IV. Outlook: From Project Delivery to Continuous Modernization

The global cloud migration services market continues to grow rapidly, but its growth drivers have shifted: powered by AI workload deployment and deep cloud adoption for core business — rather than simple "resource relocation." In this transition, merely stacking point tools cannot bridge the three gaps of information, execution, and capability; what customers need is a long-term partner that understands their business, can hold the bottom line, and keeps evolving.

4.1 MigraQ Platform Empowerment: From "Efficiency Tool" to "the Customer's Continuously Evolving Cloud Asset"

For customers, the value of MigraQ and its accompanying AI product suite (CodeBuddy, CloudQ) extends far beyond speed gains during migration — it consolidates a set of sustainable and operable digital assets for the customer.

A unified data foundation and asset graph: Through early multi-source reconciliation and dependency graph construction, customers no longer receive scattered ledgers but structured digital assets with high confidence and traceability. This baseline can connect seamlessly into existing CMDB and operations systems.

Trusted code and infrastructure assets: The expert-reviewed Terraform IaC module library and the remediated code baseline are delivered directly to the customer. This achieves "compliance as code" for infrastructure and lays a solid foundation for subsequent agile development and multi-cloud management.

Knowledge flywheel and governance: After migration, CloudQ continuously performs automated inspection of security, cost (FinOps), and performance. Lessons from anomaly handling and specification adjustments are fed back into the knowledge base — the platform grows smarter with use, truly helping customers leap from "getting to the cloud" to "getting the most from the cloud."

4.2 The Tencent Cloud MigraQ™ Service Center: Trustworthy Delivery Built on Mechanisms and Accountability

In the AI era, while smart tools matter, the bottom line of production migration is reliability. Through rigorous human-AI collaboration mechanisms, the Tencent Cloud MigraQ™ Service Center ensures every delivery is worthy of trust:

Not blindly trusting AI — experts bear ultimate responsibility: We insist that experts remain at the controls. For complex assessments of business interruption tolerance, risk acceptance of non-equivalent mappings, and final authorization of cutover windows, the experienced Tencent Cloud expert team always has the final say — the customer's business never pays the price of AI hallucination.

Ironclad quality gates: From mandatory-blocking Policy-as-Code guardrails to the three absolute red lines of "no cutover without SPOC sign-off, no specification downgrades without FinOps review," the Service Center locks down risk exposure with rigorous engineering standards — countering uncertainty with deterministic human approval.

Comprehensive ecosystem and capability assurance: Relying on Tencent Cloud's global technical operations system and the TCAF standards framework, combining the expertise of Tencent Cloud with that of certified ecosystem partners, we provide end-to-end assurance from research to operations — ensuring assets belong to the customer and expertise remains with the customer.

4.3 Closing Remarks

Migration is not a one-time delivery — it is the critical starting point for an enterprise to plug into continuously evolving AI capability. The Tencent Cloud MigraQ™ Service Center and the MigraQ platform stand ready to be your most reliable navigator and enabler in the uncharted waters of migration. We deeply integrate "MigraQ + experts" to help you bridge the gaps — so that the endpoint of every migration becomes a new starting point for your organization's smart leap.

Secure, Smooth and Smart

Tencent Cloud MigraQ™ Service Center